

Security Server LDAP Server Administration and Use



Security Server LDAP Server Administration and Use

Note

Before using this information and the product it supports, be sure to read the general information under "Appendix G. Notices".

Acknowledgements

Some of the material contained in this document is a derivative of LDAP documentation provided with the University of Michigan LDAP reference implementation (Version 3.3). Copyright © 1992-1996, Regents of the University of Michigan, All Rights Reserved.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by NEC Systems Laboratory.

Fourth Edition, September 2002

This edition, SC24-5923-03, applies to Version 1 Release 4 of z/OS (program number 5694-A01) and Version 1 Release 4 of z/OS.e (program number 5655-G52), and to all subsequent releases until otherwise indicated in new editions.

IBM® welcomes your comments. A form for readers' comments may be provided at the back of this document, or you may address your comments to the following address:

International Business Machines Corporation Department 55JA, Mail Station P384 2455 South Road Poughkeepsie, NY 12601-5400 United States of America

FAX (United States & Canada): 1+845+432-9405

FAX (Other Countries):

Your International Access Code +1+845+432-9405

IBMLink[™] (United States customers only): IBMUSM10(MHVRCFS)

Internet e-mail: mhvrcfs@us.ibm.com

World Wide Web: http://www.ibm.com/servers/eserver/zseries/zos/webqs.html

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- · Title and order number of this document
- · Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1998, 2002. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	 										-		xii
Tables	 												. xv
Preface					_								xvi
Who should use this document													
How this document is organized													
Conventions used in this document													
Where to find more information													
Softcopy publications.													
z/OS online library.													
Accessing licensed documents on the Web Using LookAt to look up message explanations													
Summary of Changes													
Part 1. Administration											_		
Part I. Administration	 •		•	•		•	•	•		•	٠	•	. 1
Chapter 1. Introducing the z/OS LDAP server													. 3
What is a directory service?													. 3
What is LDAP?													
What kind of information can be stored in the director													
How is the information arranged?													
How is the information referenced?	 •	•		•			·	•				·	5
How is the information accessed?	 •	•		•	•		•	•	•		•	•	. 6
How is the information protected from unauthorized a													
How does LDAP work?													
What about X.500?													
What are the capabilities of the z/OS LDAP server? .													
RFCs supported by z/OS LDAP	 •	•		٠	•		٠	•	•		-	٠	. 9
Chapter 2. Planning and roadmap.													. 11
Planning directory content													. 11
LDAP server roadmap													
Chantar 2 Installing and setting up related and distance													40
Chapter 3. Installing and setting up related products													
Installing and setting up DB2 for TDBM	 -					•			٠	٠		•	. 13
Getting DB2 installed and set up for CLI and ODBC.													
Installing RACF for SDBM and native authentication .													
Installing and setting up Policy Director and SAF for z/O													
Installing System SSL													
Installing OCSF and ICSF for password encryption	 												. 16
OCSF	 												. 16
ICSF	 												. 16
Installing Kerberos													
Objection 4. Configuration on LDAD common units of the late	 4 ! 1												4-
Chapter 4. Configuring an LDAP server using the Ida													
Overview of the LDAP configuration utility													
Capabilities													
Restrictions													
Using the Idapcnf utility													
Format													. 20
Example	 												. 20
•													

Input file description Usage Configuration roles and responsibilities Steps for configuring an LDAP server Configuration confirmation Specifying advanced configuration options with the Idapcnf utility	20 22 24 27 27
LDAP server configuration for other z/OS components or products	29
Chapter 5. Configuring an LDAP server without the Idapcnf utility	
LDAP server configuration roadmap	31
Preparing for configuration variable interactions	33
Chapter 6. Setting up the user ID and security for the LDAP server	35
Setting up a user ID for your LDAP server	
Requirements for a user ID that runs the LDAP server	
Protecting the environment for the LDAP server	
Protecting the environment for SSL/TLS	37
Protecting the environment for password encryption	
	-
Chapter 7. Preparing the backends, SSL/TLS, and password encryption	39
Copying the configuration files	39
Creating a DB2 database for the sample server	
Creating the DB2 database and table spaces for TDBM	40
Partitioning DB2 tables for TDBM	41
Setting up for SDBM	42
Running SDBM with other backends	43
Setting up for extended operations	43
Setting up for SSL/TLS	
Using SSL/TLS protected communications	
Enabling SSL/TLS support	
Creating and using a key database or key ring.	45
Obtaining a certificate	45
Setting up the security options for the LDAP server	45
Setting up an LDAP client	47
Using LDAP client APIs to access LDAP using SSL/TLS	
Support of certificate bind	
Configuring for user password encryption	47
Chapter 8. Customizing the LDAP server configuration	
Creating the slapd.conf file	
Locating slapd.conf	
Configuration file format	
Configuration file checklist	
Configuration file options	
Specifying a value for filename	
Deprecated options	
Configuration considerations	
Determining operational mode	
Operating in single-server mode	
Operating in multi-server mode without dynamic workload management enabled	
Operating in multi-server mode with dynamic workload management enabled	
Operating in PC callable support mode	
Setting up multiple LDAP servers	
Establishing the administrator DN and password	
Example configuration scenarios	82 82
AND THE RESERVE OF THE PROPERTY OF THE PROPERT	. 1/

	Configuring SDBM and TDBM backends	. 84
	Configuring and using multiple concurrent servers in a sysplex	. 85
	Chapter 9. Running the LDAP server	. 93
	Setting up the PDS for the LDAP server DLLs	
	Setting up and running the LDAP server as a started task	
	Defining the started task for the LDAP server	
	Running the LDAP server using the sample JCL	
	Started task messages	
	Running the LDAP server using data sets	
	Setting up and running the LDAP server in the z/OS shell	
	Dynamic debugging	
	Gathering trace records into the in-storage trace table	
	Customizing the trace table	
	Printing the trace table	
	Resetting the trace table	
	Interaction between debug and in-storage trace	
	Capturing performance information	
	Activity Logging.	
	Using the LDAP server for PC callable support	
	Verifying the LDAP server	102
	Chapter 40. Migrating to a 7/00 LDAD conver	400
	Chapter 10. Migrating to a z/OS LDAP server	
	Steps for migrating	
	Actions required for migrations	
	Obtaining the DB_VERSION setting for RDBM	
	Obtaining the DB_VERSION setting for TDBM	
	Coexistence and migration with previous releases (RDBM)	105
	Migrating existing RDBM data to a TDBM database	
	Migrating RDBM to TDBM: entryOwner and aclEntry attribute value changes	
	Migrating RDBM to TDBM: timestamp changes	
	Migrating from port, securePort, and security to listen option	
	Migrating from slapd command-line arguments -p and -s to -l	
	Migrating from maxThreads and waitingThreads to commThreads option	
	Migrating to the updated maxConnections option	
	APF authorizations	
	Schema migration	
	Migrating TDBM databases	
	Removing schema file include statements	
	Removing verifySchema configuration option	
	Migration roadmap	
	z/OS V1R4 update summary	
	z/OS V1R2 update summary	
	z/OS V1R1 update summary	
	OS/390 V2R10 update summary	
	z/OS V1R4 overview	
	Modify DN	
	CRAM-MD5 and DIGEST-MD5 authentication	
	Transport Layer Security (TLS) support	
ļ	ACL enhancements	
	SDBM enhancements	
ļ	LDAP IBM-entryuuid support	
	Activity logging	
	Abandon support	
- 1	RDBM removal	116

Monitor support		
TDBM schema		
LDAP server message enhancements		117
z/OS V1R3 overview		117
z/OS V1R2 overview		117
Kerberos authentication		
Native authentication		
SDBM enhancements		
Extended operations for accessing Policy Director data		
APF authorization		
listen configuration option		
commThreads migration		
maxConnections migration		
z/OS V1R1 overview		
OS/390 V2R10 overview		
TDBM backend		121
Schema publication		121
Schema update		122
LDAP server configuration utility		
Summary of interface changes		
Chapter 11. Running and using the LDAP backend utilities		125
Running the LDAP TDBM backend utilities in the z/OS shell		
Running the LDAP TDBM backend utilities from JCL		
Running the LDAP TDBM backend utilities in TSO		
Idif2tdbm program		
tdbm2ldif program	•	131
Charter 40. Burning and using the LDAD auto-1111D utility.		
Chapter 12. Running and using the LDAP entry UUID utility		
Running the Idapadduuids utility from JCL		
Running the Idapadduuids utility in the z/OS shell		
Idapadduuids utility		
SSL/TLS information for LDAP utilities		
Using RACF key rings		145
Chapter 13. Running and using the password encryption utility		147
Running the db2pwden utility from JCL		147
Running the db2pwden utility in the z/OS shell		147
db2pwden utility		148
Chapter 14. Internationalization support	_	151
Translated messages		
UTF-8 support		
	•	101
Part 2. Use	_	52
Part 2. USe	- 1	55
OL 4 45 D 4 11		
Chapter 15. Data model		
Relative distinguished names		
Distinguished name syntax		
Domain component naming		
RACF-style distinguished names		156
Chapter 16. LDAP directory schema for TDBM		
Setting up the schema for TDBM - new users		159
Upgrading schema for TDBM		
Updates to the schema		

	Schema introduction				100
	Schema attribute syntax				166
	LDAP schema attributes				
	Defining new schema elements				172
	Updating the schema				
	Analyzing schema errors				
	Retrieving the schema				
	Displaying the schema entry				
	Finding the subschemaSubentry DN				
	Migrating the schema from RDBM to TDBM				
ı	IBM supplied schema				
l I	User-defined schema				
l	The schema2ldif utility				
	The Schemazian unity	•	•	•	177
	Chapter 17. Modify DN Operations				102
l I	Modify DN Operation Syntax				
 	Considerations in the use of Modify DN operations.				
	Eligibility of entries for rename				
	Concurrency considerations between Modify DN operations and other LDAP operations				
	Access control and ownership				
	Relocating an entry				
	Relocating an entry with DN realignment requested				
	Access control changes				
	Ownership changes				
	Modify DN operations related to suffix DNs				
	Scenario constraints				
	Example scenarios				194
	Modify DN operations and replication				200
	Periodic validation of compatible server versions in replica servers				201
	Loss of replication synchronization due to incompatible replica server versions				
 	Loss of replication synchronization due to incompatible replica server versions Loss of replication synchronization due to incompatible replica server versions - recovery				201
 					201 202
 	Loss of replication synchronization due to incompatible replica server versions - recovery				201 202 202
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 202
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations			 	201 202 202 203 205
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility			 	201 202 202 203 205
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations			 	201 202 202 203 205 206
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations			 	201 202 202 203 205 206 210
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations			 	201 202 202 203 205 206 210 211
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations			 	201 202 203 203 205 210 211 211
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations				201 202 203 203 205 210 211 211 212
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations				201 202 203 203 205 210 211 212 216
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations				201 202 203 203 205 210 211 212 216 217
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations				201 202 203 206 206 210 211 211 212 216 217 219
 	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 205 206 210 211 211 212 216 217 218 220
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility Activity log and console listing of Modify DN operations				201 202 203 205 206 210 211 211 212 216 217 218 220
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility. Activity log and console listing of Modify DN operations				201 202 203 205 206 210 211 211 212 216 217 219 220 222
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 203 206 210 211 211 212 216 227 222 222
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility. Activity log and console listing of Modify DN operations. Chapter 18. Accessing RACF information. Mapping LDAP-style names to RACF attributes. RACF namespace entries. SDBM schema information. SDBM support for pound sign. SDBM operational behavior. Mapping SDBM requests to RACF commands. SDBM search capabilities. Using SDBM to change a user password in RACF. Using LDAP operation utilities with SDBM. Deleting attributes. Chapter 19. Kerberos authentication. Setting up for Kerberos.				201 202 203 203 206 210 211 211 212 216 227 222 222 225 225
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility. Activity log and console listing of Modify DN operations. Chapter 18. Accessing RACF information Mapping LDAP-style names to RACF attributes. RACF namespace entries. SDBM schema information. SDBM support for pound sign. SDBM operational behavior. Mapping SDBM requests to RACF commands. SDBM search capabilities. Using SDBM to change a user password in RACF. Using LDAP operation utilities with SDBM. Deleting attributes. Chapter 19. Kerberos authentication Setting up for Kerberos. Updating the directory schema for Kerberos.				201 202 203 203 206 210 211 212 216 217 222 222 225 225 226
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 206 210 211 212 216 217 219 220 222 225 226 226
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 205 206 210 211 212 216 217 219 220 222 226 226 227
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 203 206 210 211 211 212 216 227 222 226 227 227 227
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 203 206 210 211 211 212 216 227 222 226 227 227 227 227
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 203 206 210 211 211 212 216 227 222 227 227 227 227 227
	Loss of replication synchronization due to incompatible replica server versions - recovery Replication of Modify DN operations, shared databases, and compatibility				201 202 203 203 206 210 211 211 212 216 217 222 222 227 227 227 227 227 227 227

Chapter 20. Native authentication	. 233
Initializing native authentication	. 233
Updating the schema for native authentication	
Defining participation in native authentication	
Updating native passwords	
Example of setting up native authentication	
Using native authentication with Web servers	. 241
Chapter 21. CRAM-MD5 and DIGEST-MD5 Authentication	
Considerations for setting up a TDBM backend for CRAM-MD5 and DIGEST-MD5 Authentication	
CRAM-MD5 and DIGEST-MD5 configuration parameter	
Example of setting up for CRAM-MD5 and DIGEST-MD5	. 244
Chapter 22. Using extended operations to access Policy Director data.	. 247
GetDnForUserid extended operation	
GetPrivileges extended operation	
Chapter 23. Using access control	240
Access control attributes	
aclEntry attribute	
aclPropagate attribute	
aclSource attribute	
entryOwner attribute	. 253
ownerPropagate attribute	. 254
ownerSource attribute	. 254
Initializing ACLs with TDBM	
Access determination	
Attribute classes and searching	
Filter	
Compare	
Requested attributes	
Propagating ACLs	
Example of propagation	. 257
Examples of overrides	. 257
Other examples	
Access control groups	
Retrieving ACL information from the LDAP server	
Creating and managing access controls	
Creating an ACL	
Modifying an ACL	
Deleting an ACL	
Creating an owner for an entry	
Modifying an owner for an entry	. 264
Deleting an owner for an entry	. 266
Creating a group for use in ACLs and entry owner settings	. 266
Modifying a group	
Deleting a group	
	. 200
Chanter 24 Poplication	260
Chapter 24. Replication	
Ibm-entryuuid replication	
Complex modify DN replication	
Password encryption and replication	
Benefits of replication	. 270
Master server	. 270
Replica objects	
Localhost suffix	

Adding replica objects in TDBM																			. 272
Replica server																			
Populating a replica									Ċ										. 272
Configuring the replica																			
LDAP update operations on replicas .																			
Changing a replica to a master																			
SSL/TLS and replication																			
Replica server with SSL/TLS enablemen																			
Master server with SSL/TLS enablement																			
Troubleshooting																			
Recovering from out-of-sync conditions																			
Recovering from out-or-sync conditions		•		•	•	•		 •	•	•	•	•	•	•	•	•	•	•	. 213
Chapter 25. Referrals																			277
Using the referral object class and the ref																			
Creating entries																			
Associating servers with referrals																			
Pointing to other servers																			
Defining the default referral																			
Processing referrals																			
Using LDAP Version 2 referrals																			
Using LDAP Version 3 referrals																			
Example: associating servers through refer	rals	anc	rep	lica	tioi	n		 ٠	٠				•		٠	٠	٠		. 280
Chantar 2C Organizing the directory no			_																207
Chapter 26. Organizing the directory nar	nesp	Jac	е.	•	•	•	•	 •	•	•	•	•	•	•	•	•	•	•	. 207
Information layout																			
Example of building an enterprise directory																			
Priming the directory servers with informati																			
Using LDIF format to represent LDAP er																			
Generating the file																			
Setting up for replication																			
Defining another LDAP server																			
Preparing the replica																			
Notifying users of the replica																			. 294
Chapter 27. Client considerations																			
Root DSE																			
Monitor Support																			
CRAM-MD5 Authentication Support																			
UTF-8 data over the LDAP Version 2 proto																			
Attribute types stored and returned in lowe	rcase	Э																	. 298
Abandon behavior																			. 298
Changed reason codes																			. 298
Reason codes																			. 299
Part 3. Messages						٠.								٠.				٠.	315
_																			
Chapter 28. LDAP server messages																			. 317
LDAP server messages (0000)																			
Administration messages (2000)																			
TDBM messages (3000)																			
LDAP server product messages (4000) .																			
SDBM messages (5000)																			
Idapcnf messages (7000)																			
Extended operations messages (9000) .																			
	-	-		-	-	-		-		-	-	-	-	-	-	-	-	-	

Part 4. Appendixes	69
Appendix A. LDAP server configuration file (slapd.conf)	
Appendix B. Minimum schema for TDBM	393
Appendix C. SDBM schema	397
Appendix D. SPUFI files	403
Appendix E. Sample JCL.4Sample JCL for the LDAP server4Sample JCL for Idif2tdbm4Sample JCL for tdbm2ldif4Sample JCL for Idapadduuids4	413 415 417
Appendix F. Sample LDIF input file	121
Appendix G. Supported server controls4IBMModifyDNTimelimitControl4IBMModifyDNRealignDNAttributesControl4manageDsalT4authenticateOnly4IbmLDAPProxyControl4	431 431 432 432
Appendix H. Supported extended operations 4 GetDnForUserid 4 GetPrivileges 4 Start TLS 4	435 436
Appendix I. TDBM schema migration	439 439
Migrating TDBM schema from a z/OS V1R2 or z/OS V1R3 LDAP server to a z/OS V1R4 LDAP server	142
Appendix J. Accessibility	147
Notices	
Bibliography	453 453 453 453

Glossary	٠.																		455
Index .																			459

Figures

		Directory hierarchy example
	2.	Sample DSNAOINI file
	3.	LDAP configuration utility overview
		Sample portion of Idap.profile
ı		LDAP configuration utility roles and responsibilities
'		General format of slapd.conf
	7.	Multi-server sample configuration (phase 1)
	8.	Configuration file for Server A on hosta
		Contents of ABCCO.DB2CLI.CLIINIA
		Configuration file for Server B on hostb
		Contents of ABCCO.DB2CLI.CLIINIB
		Configuration file for Server C on hostc
		Contents of ABCCO.DB2CLI.CLIINIC
		Multi-server sample configuration (phase 2)
		Multi-server sample configuration (phase 3)
		Sample Schema Entry
		Object class hierarchy example
		Example attribute type file (at.in)
		Example object class file (oc.in)
		Before Modify DN operation
	21.	After Modify DN operation
		Before Modify DN operation
	23.	After Modify DN operation
	24.	Before Modify DN operation
ĺ	25.	After Modify DN operation
İ		Before Modify Dn operation
İ		After Modify DN operation
İ		Suffix rename with no new superior
' 		Suffix rename with new superior
l I		Overlapping suffix rename A
l I		Overlapping suffix rename B
l I		Suffix rename to non-suffix entry
l I		Rename non-suffix entry to suffix entry
l		RACF namespace hierarchy
		Kerberos directory example
		Native authentication example
		Example of adding propagating ACL to existing entry in directory
	٥1.	Example of adding propagating ACL to existing entry in directory
		Example of adding propagating ACL to existing entry in the directory
		Example of setting up a non-propagating ACL
	40.	Example of adding an aclEntry attribute value
	41.	Example of modifying aclPropagate attribute
	42.	Example of removing a single aclEntry attribute value
	43.	Example of deleting an ACL from an entry
	44.	Example of adding a propagating set of entry owners to existing entry in the directory 26
	45.	Example of setting up a non-propagating entry owner
	46.	Example of adding an entryOwner attribute value
	47.	Example of modifying the ownerPropagate attribute
	48.	Example of removing a single entryOwner Attribute value
	49.	Example of deleting an entry owner set from an entry
	50.	Example of creating an accessGroup entry in the directory
	51.	Example of adding a group to access control information
	52.	Example of adding a group to entry owner information
	53.	Adding a member to an accessGroup
		•

54.	Deleting a member from an accessGroup
55.	Deleting an entire accessGroup entry from the directory
56.	Example using ref attribute
57.	Setting up the servers
58.	Server A database (LDIF input)
59.	Server D configuration file
60.	Referral example summary (servers A and E)
61.	Referral example summary (servers B1 and B2)
62.	Referral example summary (servers C and D)
63.	Chicago base configuration
64.	Los Angeles base configuration
65.	New York City base configuration

Tables

	1.	LDAP server roadmap	11
	2.	LDAP configuration utility roles and responsibilities	23
		LDAP server configuration roadmap	
		Configuration variable interactions	
		TDBM value overview	
		TDBM table space partitioning indexes and values	
		TDBM table space partitioning using EID range	
		SDBM with other backends	
		Configuration file options checklist	
1		Supported ciphers	
'		Configuration considerations	
		Sample checklist and slapd.conf (using TDBM, SSL/TLS, and password encryption)	
		Sample checklist and slapd.conf (using SDBM and TDBM)	
		Sample checklist and slapd.conf (using EXOP)	
		Sample checklist (using TDBM with multi-server and dynamic workload management)	
		Debug levels	
ı		Summary of LDAP server updates for z/OS V1 R4	
1		Summary of LDAP server updates for z/OS V1R2	
		Summary of LDAP server updates for OS/390 V2R10	
		Considerations for using Idif2tdbm or Idapadd	
		db2pwden options	
		Mapping between Unicode and UTF-8	
		Syntax and default EQUALITY matching rule	
		Syntax and acceptable matching rules (EQUALITY, ORDERING, and SUBSTR)	
		Character representations	
		Supported LDAP syntaxes - general use.	
		Supported LDAP syntaxes - server use	
		Supported matching rules	
		The errno values returned by _passwd()	
		Mapping of LDAP-style names to RACF attributes (user).	
		Mapping of LDAP-style names to RACF attributes (group)	
		Mapping of LDAP-style names to RACF attributes (connection)	
		RACF backend behavior	
		SDBM-supported search filters	
		RACF backend search filters	
		Kerberos attributes and object classes	
		Operating modes for native authentication	
ı		The errno values returned by _passwd()	
'		Behavior of native authentication in example 1	
	40.		
ı	41.		
1	42.	TDBM ACL and entry owner attributes	
	43.	Permissions which apply to an entire entry	
	44.	Permissions which apply to attribue access classes	
	45.	Replica object-schema definition (mandatory attributes)	
	46.		
		-L/	

Preface

This document supports z/OS (5694-A01) and z/OS.e (5655-G52) and explains the LDAP server. The LDAP server supports Lightweight Directory Access Protocol (LDAP) and runs as a stand-alone daemon. It is based on a client/server model that provides client access to an LDAP server. The LDAP server provides an easy way to maintain directory information in a central location for storage, updating, retrieval, and exchange.

Who should use this document

This document is intended to assist system administrators. System administrators should be experienced and have previous knowledge of directory services. It is also intended for anyone that will be implementing the directory service.

How this document is organized

This document is divided into the following parts:

- Part 1, "Administration" on page 1
- Part 2, "Use" on page 153
- Part 3, "Messages" on page 315
- "Part 4. Appendixes," on page 371

Conventions used in this document

This document uses the following typographic conventions:

Bold words or characters represent API names, attributes, status codes, environment variables, parameter values, and system elements that you must enter into the system literally, such commands, options, or path names.

Italic Italic words or characters represent values for variables that you must supply.

Example Font

Examples and information displayed by the system appear in constant width type style.

- [] Brackets enclose optional items in format and syntax descriptions.
- Braces enclose a list from which you must choose an item in format and syntax descriptions.
- A vertical bar separates items in a list of choices.
- Angle brackets enclose the name of a key on the keyboard.
- ... Horizontal ellipsis points indicate that you may repeat the preceding item one or more times.
- A backslash is used as a continuation character when entering commands from the shell that exceed one line (255 characters). If the command exceeds one line, use the backslash character \ as the last nonblank character on the line to be continued, and continue the command on the next line.

Where to find more information

Where necessary, this document references information in other documents. For complete titles and order numbers of the documents for all products that are part of z/OS, refer to z/OS: Information Roadmap.

For a list of titles and order numbers of the documents that are useful for z/OS LDAP, see "Bibliography" on page 453.

Softcopy publications

The z/OS Security Server library is available on a CD-ROM, z/OS: Collection, SK3T-4269. The CD-ROM online library collection is a set of unlicensed books for z/OS and related products that includes the IBM Library Reader. This is a program that enables you to view the BookManager files. This CD-ROM also contains the Portable Document Format (PDF) files. You can view or print these files with the Adobe Acrobat reader.

z/OS online library

The softcopy z/OS publications are also available for web browsing and for viewing or printing PDFs using the following URL:

http://www.ibm.com/servers/eserver/zseries/zos/bkserv

You can also provide comments about this document and any other z/OS documentation by visiting that URL. Your feedback is important in helping to provide the most accurate and high-quality information.

Accessing licensed documents on the Web

z/OS licensed documentation in PDF format is available on the Internet at the IBM Resource Link Web site:

http://www.ibm.com/servers/resourcelink

Licensed documents are available only to customers with a z/OS license. Access to these documents requires an IBM Resource Link user ID, password, and z/OS licensed book key code. The z/OS order you received provides a memo that includes your key code.

To obtain your IBM Resource Link user ID and password, logon to:

http://www.ibm.com/servers/resourcelink

To register for access to the z/OS licensed books:

- Logon to Resource Link using your Resource Link user ID and password.
- 2. Select **User Profiles** located on the left-hand navigation bar.
- 3. Select Access Profile.
- 4. Select Request Access to Licensed books.
- 5. Supply your key code where requested and select the **Submit** button.

If you supplied the correct key code you will receive confirmation that your request is being processed.

After your request is processed you will receive an e-mail confirmation.

Note: You cannot access the z/OS licensed books unless you have registered for access to them and received an e-mail confirmation informing you that your request has been processed.

To access the licensed books:

- 1. Log on to Resource Link using your Resource Link user ID and password.
- 2. Select Library.
- 3. Select zSeries.
- 4. Select Software.
- 5. Select z/OS.
- 6. Access the licensed book by selecting the appropriate element.

Using LookAt to look up message explanations

LookAt is an online facility that allows you to look up explanations for z/OS messages, system abends, and some codes. Using LookAt to find information is faster than a conventional search because in most cases LookAt goes directly to the message explanation.

You can access LookAt from the Internet at:

http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookat.html

or from anywhere in z/OS where you can access a TSO command line (for example, TSO prompt, ISPF, z/OS UNIX System Services running OMVS).

To find a message explanation on the Internet, go to the LookAt Web site and simply enter the message identifier (for example, IAT1836 or IAT*). You can select a specific release to narrow your search. You can also download code from the *z/OS: Collection* and the LookAt Web site so you can access LookAt from a PalmPilot (Palm VIIx suggested).

To use LookAt as a TSO command, you must have LookAt installed on your host system. You can obtain the LookAt code for TSO from a disk on your *z/OS: Collection* or from the LookAt Web site. To obtain the code from the LookAt Web site, do the following:

- 1. Go to
 - http://www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookat.html
- 2. Click the News button.
- 3. Scroll to Download LookAt Code for TSO and VM.
- 4. Click the ftp link, which will take you to a list of operating systems. Select the appropriate operating system. Then select the appropriate release.
- 5. Find the **lookat.me** file and follow its detailed instructions.

To find a message explanation from a TSO command line, simply enter: **lookat** *message-id*. LookAt will display the message explanation for the message requested.

Note: Some messages have information in more than one book. For example, IEC192I has routing and descriptor codes listed in *z/OS: MVS Routing and Descriptor Codes*. For such messages, LookAt prompts you to choose which book to open.

Summary of Changes

Summary of changes for SC24-5923-03 z/OS Version 1 Release 4

This document contains information previously presented in *z/OS Security Server LDAP Server Administration and Use*, SC24-5923-02, which supports z/OS Version 1 Release 2 and z/OS Version 1 Release 3.

The following summarizes the changes to that information:

New information

- · CRAM-MD5 (RFC 2195) and DIGEST-MD5 (RFC 2831) authentication
- · New global CRAM-MD5 and DIGEST-MD5 configuration option: digestRealm
- · LDAP IBM-entryuuid support
- New LDAP entry UUID configuration option: serverEtherAddr
- New LDAP entry UUID utility: Idapadduuids
- The LDAP server supports both Secure Sockets Layer (SSL) Version 3 and Transport Layer Security
 (TLS) Version 1 for secure protected sessions. Where the book discusses both SSL and TLS together,
 they are written as SSL/TLS. Support for client connections that begin on the non-secure port which
 then switch to a secure (SSL/TLS) protected session using the LDAP Start TLS extended operation.
 This support implements RFC 2830.
- ACL enhancements to allow attribute-level access control and the ability to explicitly deny access to information.
- Full support for Modify DN operations, including subtree relocation
- New supported server controls: IBMModifyDNTimelimitControl and IBMModifyDNRealignDNAttributesControl
- · Server activity logging
- · New activity logging configuration option: logfile
- Monitor support
- Migration information
- New LDAP server messages and updated severity levels new messages:
 - GLD0207I GLD0241A
 - GLD2100A GLD2125A
 - GLD3133A GLD3143I
 - GLD4013A GLD4020I
 - GLD5006A
- Information is added to indicate this document supports z/OS.e

Changed information

- The **db2pwden** utility requires that the value specified with the **-N** operand be the certificate label. The certificate name for this operand is no longer supported.
- The following configuration options are now ignored by the LDAP server: attribute, index, objectclass, verifySchema, tbspaceentry, tbspacemutex, tbspace32k, and tbspace4k.
- With the removal of the RDBM backend, **rdbm** can no longer be specified as the *dbtype* on the **database** configuration option.
- The **database** configuration option has been extended to allow specification of a unique name for each configured backend instance.

- The slapd.conf file has been updated.
- Severity indicators have changed for most server messages.
- The instructions for migrating to new levels of TDBM schema have been changed.

Removed information

- RDBM backend (including the Idif2db and db2Idif utilities and the Idapcp command)
- Online documentation in HTML form for the z/OS LDAP client APIs. Prior to s/OS V1T4, the LDAP Programming guide documentation was shipped in html form as part of the LDAP product. This documentation is no longer shipped. Refer to "z/OS online library" on page xviii for information about accessing online copies of this documentation.

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

Summary of changes for SC24-5923-02 z/OS Version 1 Release 2 as updated December 2001

This book contains information previously presented in z/OS SecureWay Security Server LDAP Server Administration and Use, SC24-5923-01, which supports z/OS Version 1 Release 2.

The following summarizes the changes to that information:

New information

- Program Call (PC) callable support for Policy Director, including a new global configuration option, pcThreads
- New extended operations that retrieve Policy Director data: GetDnForUserid and GetPrivileges. Support for these operations is provided by a new extended operations (EXOP) backend.
- New supported server control: IbmLDAPProxyControl

This book contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

Summary of changes for SC24-5923-01 z/OS Version 1 Release 2

This book contains information previously presented in z/OS SecureWay Security Server LDAP Server Administration and Use, SC24-5923-00, which supports z/OS Version 1 Release 1.

The following summarizes the changes to that information:

New information

- Kerberos authentication
- New global Kerberos configuration options: serverKrbPrinc, krbKeytab, krbLDAPAdmin, supportKrb5
- New backend Kerberos configuration option: krbldentityMap
- Native authentication
- · New native authentication configuration options: useNativeAuth, nativeAuthSubtree, and nativeUpdateAllowed
- RACF connect profile support and other SDBM enhancements

- Additional z/OS LDAP schema files
- Changes to APF-authorization
- · New threading-related configuration option: commThreads
- New front-end configuration options: idleConnectionTimeout and listen
- Debug level specification
- Migration information
- New LDAP server messages GLD0162I GLD0206E, GLD3124E GLD3132E, GLD5004E, and **GLD5005E**

Changed information

- The slapd.conf file has been updated.
- · The minimum schema for TDBM has been updated.
- Use of the following configuration options is deprecated: waitingThreads, maxThreads, security, port, and securePort. The waitingThreads and maxThreads options are replaced by commThreads. The security, port, and securePort are replaced by listen.
- Use of the following slapd startup options is deprecated: -s and -p. These are replaced by -l.

This book contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

You may notice changes in the style and structure of some content in this book - for example, headings that use uppercase for the first letter of initial words only. The changes are ongoing improvements to the consistency and retrievability of information in our books.

Summary of changes for SC24-5923-00 z/OS Version 1 Release 1

This following highlights function that is documented in this book:

- An LDAP configuration utility, Idapcnf, that simplifies and automates the LDAP server configuration process for TDBM and, optionally, SDBM.
- New LDAP server messages GLD7001I-GLD7011I.

Part 1. Administration

Chapter 1. Introducing the z/OS LDAP server

The z/OS Lightweight Directory Access Protocol (LDAP) server, part of the Security Server for z/OS, is based on a client/server model that provides client access to an LDAP server. An LDAP directory provides an easy way to maintain directory information in a central location for storage, update, retrieval, and exchange.

The LDAP server provides the following functions:

- · Interoperability with other LDAP clients
- · Access controls on directory information
- Secure Sockets Layer (SSL) communication (SSL V3 and TLS V1)
- Start TLS activation of secure communication
- Client and server authentication using SSL/TLS
- Password encryption
- Replication
- Referrals
- LDAP Version 2 and Version 3 protocol support
- · Schema publication and update
- · Kerberos authentication
- · Native authentication
- CRAM-MD5 (Challenge-Response Authentication Method) and DIGEST-MD5 authentication
 - Root DSE information
 - LDAP access to information stored in RACF[®]
 - Use of DB2[®] data sharing in sysplex configurations

This book describes how to install, configure, and run the stand-alone LDAP server and other LDAP programs. It is intended for newcomers and experienced administrators alike. This section provides a basic introduction to directory services, and the directory service provided by the LDAP server in particular.

z/OS: Security Server LDAP Client Programming describes the LDAP client application programming interfaces (APIs) you can use to develop LDAP applications.

What is a directory service?

A directory is like a database, but tends to contain more descriptive, attribute-based information. The information in a directory is generally read much more often than it is written. As a consequence, directories do not usually implement the complicated transaction or rollback schemes that relational databases use for doing high-volume complex updates. Directory updates are typically simple all-or-nothing changes, if they are allowed at all. Directories are tuned to give quick-response to high-volume lookup or search operations. They may have the ability to replicate information widely in order to increase availability and reliability, while reducing response time. When directory information is replicated, temporary inconsistencies between the replicas are considered acceptable, as long as they get in sync eventually.

There are many different ways to provide a directory service. Different methods allow different kinds of information to be stored in the directory, place different requirements on how that information can be referenced, queried and updated, how it is protected from unauthorized access, and so on. Some directory services are local, providing service to a restricted context (for example, the finger service on a single machine). Other services are global, providing service to a much broader context (for example, the entire Internet). Global services are usually distributed, meaning that the data they contain is spread across many machines, all of which cooperate to provide the directory service. Typically a global service defines a uniform namespace which gives the same view of the data no matter where you are in relation to the data itself.

What is LDAP?

The LDAP server's model for the directory service is based on a global directory model called LDAP, which stands for the Lightweight Directory Access Protocol. LDAP Version 2 (V2) and LDAP Version 3 (V3), both supported in z/OS, are directory service protocols that run over TCP/IP. The details of LDAP V2 are defined in Internet Engineering Task Force (IETF) Request for Comments (RFC) 1777 *The Lightweight Directory Access Protocol* and the details of LDAP V3 are defined in the set of IETF RFCs 2251 - 2256. "RFCs supported by z/OS LDAP" on page 9 shows the entire list of supported RFCs.

This section gives an overview of LDAP from a user's perspective.

What kind of information can be stored in the directory?

The LDAP directory service model is based on *entries*. An entry is a collection of attributes that has a name, called a *distinguished name* (DN). The DN is used to refer to the entry unambiguously. Each of the entry's attributes has a type and one or more values. The types are typically mnemonic strings, like **cn** for common name, or **mail** for e-mail address. The values depend on what type of attribute it is. For example, a **mail** attribute might contain an e-mail address with an attribute value of thj@vnet.ibm.com. A **jpegPhoto** attribute would contain a photograph in binary JPEG format.

How is the information arranged?

In LDAP, directory entries are arranged in a hierarchical tree-like structure that sometimes reflects political, geographic or organizational boundaries. Entries representing countries appear at the top of the tree. Below them are entries representing states or national organizations. Below them might be entries representing people, organizational units, printers, documents, or just about anything else you can think of. Figure 1 on page 5 shows an example LDAP directory tree, which should help make things clear.

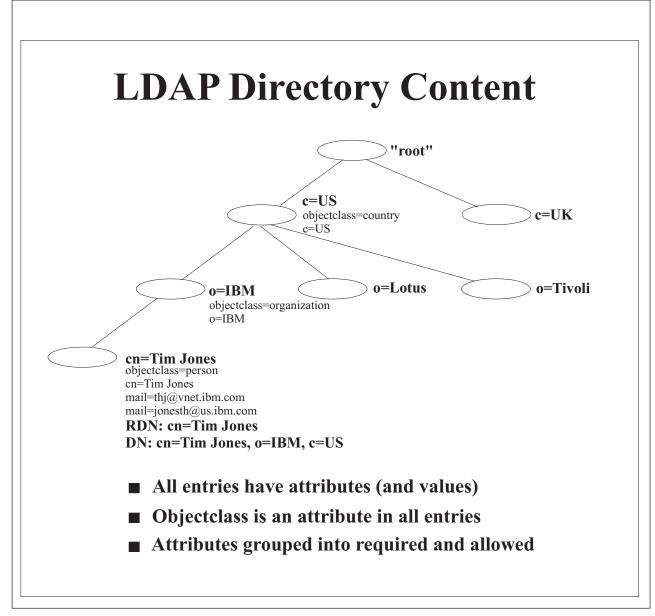


Figure 1. Directory hierarchy example

In addition, LDAP allows you to control which attributes are required and allowed in an entry through the use of a special attribute called *object class*. The values of the **objectClass** attribute determine the attributes that can be specified in the entry.

How is the information referenced?

An entry is referenced by its distinguished name, which is constructed by taking the name of the entry itself (called the *relative distinguished name*, or RDN) and concatenating the names of its ancestor entries. For example, the entry for Tim Jones in the example above has an RDN of cn=Tim Jones and a DN of cn=Tim Jones, o=IBM, c=US. The full DN format is described in IETF RFC 2253, *LDAP (V3): UTF-8 String Representation of Distinguished Names*.

The z/OS LDAP server supports different naming formats. While naming based on country, organization, and organizational unit is one method, another method is to name entries based on an organization's registered DNS domain name. Names of this form look like: cn=Tim Smith,dc=vnet,dc=ibm,dc=com. These naming formats can be mixed as well, for example: cn=Tim Brown,ou=Sales,dc=ibm,dc=com.

How is the information accessed?

LDAP defines operations for interrogating and updating the directory. Operations are provided for adding/deleting an entry to/from the directory, changing an existing entry, and changing the name of an entry. Most of the time, though, LDAP is used to search for information in the directory. The LDAP search operation allows some portion of the directory to be searched for entries that match some criteria specified by a search filter. Information can be requested from each entry that matches the criteria. The LDAP compare operation allows a value to be tested in an entry without returning that value to the client.

An example of search is, you might want to search the entire directory subtree below IBM for people with the name Tim Jones, retrieving the e-mail address of each entry found. LDAP lets you do this easily. Or you might want to search the entries directly below the c=US entry for organizations with the string Acme in their name, and that have a FAX number. LDAP lets you do this too. The section "How does LDAP work?" describes in more detail what you can do with LDAP and how it might be useful to you.

The LDAP bind operation is used to indicate to the LDAP server who is going to be making add/modify/search/compare or delete requests. The LDAP bind operation is an authentication process.

This authentication process can be used by distributed applications which need to implement some form of authentication.

How is the information protected from unauthorized access?

An Access Control List (ACL) provides a means to protect information stored in an LDAP directory. ACLs are used to restrict access to different portions of the directory, specific directory entries, or information within an entry. Access control can be specified for individual users or groups.

How does LDAP work?

LDAP directory service is based on a client/server model. One or more LDAP servers contain the data making up the LDAP directory tree. An LDAP client application connects to an LDAP server using LDAP APIs and asks it a question. The server responds with the answer, or with a pointer to where the application can get more information (typically, another LDAP server). With a properly constructed namespace, no matter which LDAP server an application connects to, it sees the same view of the directory; a name presented to one LDAP server references the same entry it would at another LDAP server. This is an important feature of a global directory service, which LDAP servers can provide.

What about X.500?

LDAP was originally developed as a front end to X.500, the OSI directory service. X.500 defines the Directory Access Protocol (DAP) for clients to use when contacting directory servers. DAP has been characterized as a heavyweight protocol that runs over a full OSI stack and requires a significant amount of computing resources to run. LDAP runs directly over TCP and provides most of the functionality of DAP at a much lower cost.

An LDAP server is meant to remove much of the burden from the server side just as LDAP itself removed much of the burden from clients. If you are already running an X.500 service and you want to continue to do so, you can probably stop reading this guide, which is all about running LDAP through an LDAP server without running X.500. If you are not running X.500, want to stop running X.500, or have no immediate plans to run X.500, read on.

What are the capabilities of the z/OS LDAP server?

You can use the z/OS LDAP server to provide a directory service of your very own. Your directory can contain just about anything you want to put in it. Some of the z/OS LDAP server's more interesting features and capabilities include:

- Multiple concurrent database instances (referred to as backends): The LDAP server can be configured to serve multiple databases at the same time. This means that a single z/OS LDAP server can respond to requests for many logically different portions of the LDAP tree. A z/OS LDAP server can be configured to provide access to RACF, as well as store application-specific information.
- Robust database: The LDAP server comes with a TDBM backend database based on DB2. The TDBM database is a highly scalable database implementation.

Note: To use TDBM, DB2 is required.

- Loading and unloading data: The LDAP server can load large numbers of entries into a TDBM DB2 database using the Idif2tdbm utility. See "Idif2tdbm program" on page 127 for more information. The LDAP server can also unload large numbers of entries from a TDBM DB2 database using the tdbm2ldif utility. See "tdbm2ldif program" on page 137 for more information.
- Access control: The LDAP server provides a rich and powerful access control facility, allowing you to control access to the information in your database or databases. You can control access to entries based on LDAP authentication information, including users and groups. Access control is configurable down to individual attributes within entries. Also, access controls can be set up to explicitly deny access to information. See Chapter 23, "Using access control" on page 249 for more information.
- Threads: The LDAP server is threaded for high performance. A single multi-threaded z/OS LDAP server process handles all incoming requests, reducing the amount of system overhead required.
- Replication: The LDAP server can be configured to maintain replica copies of its database. This master/slave replication scheme is vital in high-volume environments where a single LDAP server just does not provide the necessary availability or reliability. See Chapter 24, "Replication" on page 269 for more information. This feature is contrasted with multiple concurrent servers.
- · Referrals: The LDAP server provides the ability to refer clients to additional directory servers. Using referrals you can distribute processing overhead, distribute administration of data along organizational boundaries, and provide potential for widespread interconnection beyond an organization's own boundaries. See Chapter 25, "Referrals" on page 277 for more information.
- Configuration:

The LDAP server configuration process can be simplified by using the Idapcnf configuration utility. This utility requires minimal user interaction and allows novice LDAP users to quickly configure an LDAP server. See Chapter 4, "Configuring an LDAP server using the Idaponf utility" on page 17 for more information.

If you do not use the Idapcnf utility, the LDAP server is highly configurable through a single configuration file which allows you to change just about everything you would ever want to change. Configuration options have reasonable defaults, making your job much easier. See "Creating the slapd.conf file" on page 51 for more information.

- Secure communications: The LDAP server can be configured to encrypt data to and from LDAP clients using the z/OS Cryptographic Services System SSL. The LDAP server supports the Start TLS extended operation to switch a non-secure connection to a secure connection. It has a variety of ciphers for encryption to choose from, all of which provide server and optionally client authentication through the use of X.509 certificates. See "Setting up for SSL/TLS" on page 43 for more information.
- Multiple concurrent servers: The LDAP server can be configured to permit multiple instances to serve the same DB2-based backing store at the same time. The multiple server instances may run on the same z/OS image, and they may run on multiple z/OS images in a Parallel Sysplex®. This support is available for the TDBM backend. This improves availability and may offer improved performance in certain configurations. See "Determining operational mode" on page 75 for more information.
- **Dynamic workload management**: The LDAP server can be configured to participate in dynamic workload management in a Parallel Sysplex by exploiting TCP/IP connection optimization. With multiple concurrent server instances configured in this way, availability is improved, as is resource utilization. In addition, performance improvements may be experienced as sysplex resource utilization is more evenly balanced across z/OS systems in the sysplex. See "Determining operational mode" on page 75 for more information.

 Access to RACF data: The LDAP server can be configured to provide read/write access to RACF user, group, and connection profiles using the LDAP protocol. (RACF is a component of the Security Server for z/OS.) If the RACF data is shared across the sysplex, then users, groups, and connections in the sysplex can be managed using LDAP. The LDAP server's access to RACF is managed by an additional configurable backend called SDBM. See Chapter 18, "Accessing RACF information" on page 205 for more information.

Note: To use SDBM for ONLY authentication (LDAP bind processing), any security manager implementing the SAF service required by the __passwd() function call can be used. To use SDBM for accessing and updating USER and GROUP profile information, RACF is required.

- Retrieve Policy Director data: The z/OS LDAP server, when using the EXOP backend, supports two LDAP extended operations, GetDnForUserid and GetPrivileges, that retrieve Policy Director data from any LDAP server. See Chapter 22, "Using extended operations to access Policy Director data" on page 247 for more information.
- Native authentication: The z/OS LDAP server allows clients to bind to entries in a TDBM backend by using the system for verifying the authentication attempt. The client can perform a simple bind supplying an LDAP DN of an entry in a TDBM backend along with a security manager-maintained password. Password authentication is then performed by the security manager. See Chapter 20, "Native authentication" on page 233 for more information.

Note: To use native authentication, any security manager implementing the SAF service required by the passwd() function call can be used.

- LDAP Version 3 protocol support: The LDAP server provides support for Version 3 of the LDAP protocol. This includes:
 - All protocol operations
 - Implicit bind
 - Certificate (or Simple Authentication and Security Layer) bind
 - Version 3 referrals
 - Controls
 - Root DSE support
 - Internationalization (UTF-8) support
 - Modify name supported for all entries including subtree move
 - Schema publication (TDBM and SDBM)
 - Additional syntax support (TDBM)
 - Online schema update capability (TDBM)
- Dynamic schema: The LDAP server, when using the TDBM backend, allows the schema to be changed dynamically through the LDAP protocol. See Chapter 16, "LDAP directory schema for TDBM" on page 159 for more information.
- Internationalization (UTF-8) support: The LDAP server allows storage, update and retrieval, through LDAP operations, of national language data using LDAP Version 3 protocol. See "UTF-8 support" on page 151 for more information.
- SASL external bind and client and server authentication: The LDAP server allows client applications to use a certificate when communicating with the server using SSL/TLS communications. In order to use a certificate on bind, the server must be configured to perform both client and server authentication. This ensures both entities are who they claim to be. See "Setting up for SSL/TLS" on page 43 for more
- SASL GSS API Kerberos bind with mutual authentication: The LDAP server allows clients to bind to the server using Kerberos credentials. Mutual authentication is used to verify both the client and server identities. See Chapter 19, "Kerberos authentication" on page 225 for more information.
- SASL CRAM-MD5 and DIGEST-MD5 authentication: The LDAP server allows clients to bind to the server using DIGEST-MD5 (RFC 2831) and CRAM-MD5 (Challenge-Response Authentication Method -RFC 2195) authentication bind methods. See Chapter 21, "CRAM-MD5 and DIGEST-MD5 Authentication" on page 243 for more information.

- Support for root DSE: The LDAP server supports search operations against the Root of the Directory tree as described in IETF RFC 2251, The Lightweight Directory Access Protocol (V3). The so-called Root DSE can be accessed using LDAP V3 search operations. See "Root DSE" on page 297 and z/OS: Security Server LDAP Client Programming for more information.
- Extended group membership searching: The LDAP server supports extended group membership searching which allows the LDAP server to find a DN that may be a member of a group in a backend (TDBM) where the DN does not reside. The LDAP server can find the group memberships for the DNs in the other backends that are configured. See the **extendedGroupSearching** configuration file option on page 57 for more information.
- Supported server controls: The LDAP server supports the manageDsalT, authenticateOnly, IBMLDAPProxyControl, IBMModifyDNTimelimitControl, and IBMModifyDNRealignDNAttributesControl. See Appendix G, "Supported server controls" on page 431 for more information.
- Supported extended operations: The LDAP server supports the GetDnForUserid and GetPrivileges extended operations. See Appendix H, "Supported extended operations" on page 435
 - Password encryption: The LDAP server allows prevention of unauthorized access to user passwords stored in the TDBM backends. See "Configuring for user password encryption" on page 47 for more information.
 - Multiple socket ports: The LDAP server can be configured to listen for secure and nonsecure connections from clients on one or more interfaces on a system. With the listen configuration option on the LDAP server, the host name or the IP address or both, along with the port number, can target one or multiple interfaces on a system. Prior to z/OS Release 2 LDAP, the server would, by default, listen on all system interfaces on the nonsecure and secure ports that were configured on the LDAP server. See the **listen** configuration option on page 59 for more information.
- Ibm-entryuuid attribute: The LDAP server now generates a unique identifier for any entry that is created or modified and does not already have a unique identifier assigned. The unique identifier is stored in the **ibm-entryuuid** attribute. The **ibm-entryuuid** attribute is replicated to servers that support the **ibm-entryuuid** attribute. A utility is provided to create the **ibm-entryuuids** for existing entries when migrating from previous releases. See Chapter 12, "Running and using the LDAP entry UUID utility" on page 141 for more information on on the LDAP entry UUID utility. See the "Configuration file options" on page 53 to configure the **serverEtherAddr** keyword in the **slapd.conf** file.

Note: With APAR OW50971 for OS/390 V2R10 or z/OS V1R1, the complete LDAP product at the z/OS V1R2 level was rolled back to OS/390 V2R10 and z/OS V1R1. If this APAR is applied to an OS/390 V2R10 or z/OS V1R1 system, all functionality of the z/OS V1R2 LDAP server becomes available on OS/390 V2R10 or z/OS V1R1. If this rollback APAR has been applied to your OS/390 V2R10 or z/OS V1R1 system, you should follow sections marked with z/OS V1R2.

RFCs supported by z/OS LDAP

The z/OS LDAP server supports the following IETF RFCs:

- 1738 Uniform Resource Locators (URL)
- 1779 A String Representation of Distinguished Names
- 1823 The LDAP Application Program Interface
- 1959 An LDAP URL Format
- 1960 A String Representation of LDAP Search Filters
- 2052 A DNS RR for specifying the location of services (DNS SRV)
- 2104 HMAC: Keyed-Hashing for Message Authentication
- 2195 IMAP/POP AUTHorize Extension for Simple Challenge/Response
 - 2222 Simple Authentication and Security Layer (SASL)
 - 2247 Using Domains in LDAP/X.500 Distinguished Names
 - 2251 Lightweight Directory Access Protocol (v3)
 - 2252 Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions
 - 2253 Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names
 - 2255 The LDAP URL Format

- 2256 A Summary of the X.500(96) User Schema for use with LDAPv3
- 2279 UTF-8, a transformation format of ISO 10646
- 2743 Generic Security Service Application Program Interface Version 2, Update 1
- 2744 Generic Security Service API Version 2: C-bindings
- 2829 Authentication Methods for LDAP
- 2830 Lightweight Directory Access Protocol (v3): Extension for Transport Layer Security
- 2831 Using Digest Authentication as a SASL Mechanism
- 2849 The LDAP Data Interchange Format (LDIF) Technical Specification

Note that although the LDAP V3 protocol RFCs are listed as supported, the specific function that z/OS LDAP supports is listed in "LDAP Version 3 protocol support" on page 8.

Chapter 2. Planning and roadmap

This chapter:

- Shows you where to find information in this book that will help you plan your directory content.
- Contains a roadmap that points you to information that may be helpful in preparing for your LDAP server configuration.

Planning directory content

Before configuring and populating your database, determine:

What type of data you are going to store in the directory.

You should decide on what sort of schema you need to support the type of data you want to keep in your directory. The directory server is shipped with a standard set of attribute type and object class definitions.

Before you begin adding entries to the directory, you might need to add new attribute type and object class definitions that are customized to your data.

Schema definition styles are specific to the backend or data store being configured. Once you have determined which backends to configure, refer to Chapter 16, "LDAP directory schema for TDBM" on page 159 or "Setting up for SDBM" on page 42 for more information.

- · How you want to structure your directory data.
 - Refer to Chapter 15, "Data model" on page 155 for more information.
- A set of policies for access permissions.
 Refer to Chapter 23, "Using access control" on page 249 for more information.

LDAP server roadmap

Table 1 is a roadmap that points you to information that may be helpful in preparing for your LDAP server configuration.

See Chapter 10, "Migrating to a z/OS LDAP server" on page 103 if you have a previous release of the LDAP server installed on your system.

For complete instructions for installing the LDAP server product, see *z/OS: Program Directory* which comes with the LDAP server tape or cartridge. Be sure to read the license agreement in *z/OS: Licensed Program Specifications*, which is also included in the box.

Important

Before you proceed, review the *Memo to Users*, which describes any late changes to the procedures in this book. A printed copy is included with the LDAP server tape or cartridge.

Table 1. LDAP server roadmap

Complete?	Task	Page	
If you are config must:	If you are configuring your LDAP server for z/OS WebSphere [™] Application Server (z/OS Component Broker), you must:		
	See the WebSphere documentation for details on LDAP server requirements.		
If you are config	If you are configuring your LDAP server for z/OS Hardware Configuration Definition (HCD), you must:		
See the z/OS HCD documentation for details on LDAP server requirements.			
If you are config	If you are configuring your LDAP server for Resource Measurement Facility (RMF™), you must:		

Table 1. LDAP server roadmap (continued)

Complete?	Task	Page	
	See the z/OS RMF documentation for details on LDAP server requirements.		
If you are confi	If you are configuring your LDAP server for Managed System Infrastructure (msys), you must:		
	See the z/OS msys documentation for details on LDAP server requirements.		
If you are confi	iguring your LDAP server for z/OS Policy Director, you must:		
	See the z/OS Policy Director documentation for details on LDAP server requirements.		
If you want to	see how a working LDAP server looks, you must:		
	See the /usr/lpp/ldap/examples/sample_server directory which contains everything necessary to set up a sample LDAP server. See the README file in this directory for complete instructions.		
If you are migr	ating from a previous release of the LDAP server, you must:		
	See the migration information.	103	
If this is the firs	st time you are installing the z/OS LDAP server, you must:		
	Read the following documents that are included in the box with the z/OS LDAP server tape or cartridge: • z/OS: Program Directory, which contains the complete install instructions. • z/OS: Licensed Program Specifications, which contains the license agreement. • Memo to Users, which describes any late changes to the procedures in this book.		
Choose a conf	iguration method from the following options:		
	The Idapcnf configuration utility uses a profile file as input to generate jobs to set up the system environment and configuration. Check "Restrictions" on page 18 to decide if this method will work for your installation.	17	
	If you do not use the Idapcnf utility, use the instructions for customizing your configuration.	51	

Chapter 3. Installing and setting up related products

This chapter discusses what products are necessary to install or set up prior to configuring the z/OS LDAP server product. There are some decisions you must make depending on how you want to set up your LDAP server.

If you plan to use:	You must:	See:
TDBM backend (based on DB2)	Install the DB2 product and set up CLI and ODBC. Note that if your LDAP server will be used only for accessing RACF information, it is not necessary to install DB2 or set up a DB2 database. See "Setting up for SDBM" on page 42 for information on configuring the LDAP Server for accessing RACF information.	"Installing and setting up DB2 for TDBM" below
SDBM backend (based on RACF)	Install RACF.	"Installing RACF for SDBM and native authentication" on page 15
Program call (PC) support and the EXOP backend to support Policy Director extended operations	Install Policy Director and use SAF.	"Installing and setting up Policy Director and SAF for z/OS Policy Director support" on page 15
Protect access to your LDAP server with Secure Sockets Layer (SSL) security or Transport Layer Security (TLS)	Install z/OS Cryptographic Services System SSL.	"Installing System SSL" on page 15
Password encryption with TDBM	Install OCSF and ICSF.	"Installing OCSF and ICSF for password encryption" on page 16
Kerberos authentication	Install Kerberos.	"Installing Kerberos" on page 16
Native authentication	Install a security server.	"Installing RACF for SDBM and native authentication" on page 15

Installing and setting up DB2 for TDBM

This section describes how to get DATABASE 2^{TM} (DB2) running and how to run the LDAP server using the TDBM (DB2) backend. You should also have or have access to *DB2 for OS/390 and z/OS: ODBC Guide and Reference* and *DB2 for OS/390 and z/OS: Application Programming and SQL Guide*.

Getting DB2 installed and set up for CLI and ODBC

Following are the steps to get DB2 installed:

- 1. Have your database system administrator install DB2 for OS/390 and z/OS Version 5, Version 6, or Version 7. If you will be running your LDAP server in multi-server mode on multiple images in a Parallel Sysplex, your administrator must configure a DB2 data sharing group with members on each of the z/OS images on which an LDAP server instance will run. (See "Determining operational mode" on page 75 for a description of the various operating modes in which the LDAP server may run.)
 Make sure that the SMP/E jobs are a part of the DB2 installation. See the section about installing DB2 CLI in DB2 for OS/390 and z/OS: ODBC Guide and Reference. Also, specify the user ID (for example, suxxxx) that should be granted database system administrator authority. This should be the ID you will log on with to run the SPUFI jobs to create the DB2 tables for the LDAP server. You need to find out the following information from your database administrator:
 - DB2 subsystem name. For example, DSN5.

• DB2 server location (or data source). For example, LOC1.

In order to use a local or remote DB2 database, you must include a DDF record in your Bootstrap Data Set (BSDS). That DDF record must include a LOCATION keyword and an LUNAME keyword. If you are using a DB2 database that is on the local system (including a database that is set up for DB2 data sharing) the DDF component need not be started. If you are using a DB2 database that is on a remote system, the DDF component of DB2 must be configured and started on systems using the DB2 Call Level Interface (CLI). CLI is used by the LDAP server for requesting services from DB2. (The DB2 Call Level Interface is IBM's callable SQL interface used by the DB2 family of products, based on the ISO Call Level Interface Draft International Standard specification and the Microsoft® Open Database Connectivity specification.)

It may be necessary to have your DB2 administrator set up buffer pools, TEMP space, and TEMP datasets for additional buffer pool sizes. By default, the LDAP server DB2 backing stores will use bufferpool BP0. The bufferpool choice and size (4K, 32K or other sizes) should be examined by your database system administrator to ensure they are large enough to meet the additional needs of the LDAP server, once you have loaded data into its database. The DB2 runstats command should be used once data is loaded so that DB2 gueries are optimized.

2. Enter:

-dsn start db2

from the image console and wait for DB2 to finish the DB2 initialization. The dsn is the DB2 subsystem name.

You can stop DB2 by entering:

-dsn stop db2

from the console.

Note: This may already be done when the system is re-ipled.

- 3. Edit and Submit DSNHLQ.SDSNSAMP(DSNTIJCL) where DSNHLQ is the high-level qualifier used during DB2 installation. See the section on setting up DB2 CLI runtime environment in DB2 for OS/390 and z/OS: ODBC Guide and Reference. You must run this from the user ID that has been granted the appropriate database authorities. This step establishes the environment needed for the LDAP server to use the CLI. It is often referred to as "binding the CLI plan". When binding the CLI plan, it must be bound using the bind option RELEASE(COMMIT), either by default (when no RELEASE option is specified on the bind statement) or by explicitly specifying the option on the bind statement (see DB2 for OS/390 and z/OS: Command Reference for information on bind options and syntax). Note the plan name used when editing this job, for example DSNACLI.
- 4. Create (Allocate) DB2 CLI Initialization File. A sample of the CLI initialization file can be found at DSNHLQ.SDSNSAMP(DSNAOINI). Create your own CLI initialization file and copy DSNHLQ.SDSNSAMP(DSNAOINI) into it. If a dataset is used for the CLI initialization file, it must not contain sequence numbers. Refer to the section on the DB2 CLI initialization File in DB2 for OS/390 and z/OS: ODBC Guide and Reference for more information on the contents of this file. Figure 2 on page 15 shows a sample file. The example in Figure 2 on page 15 shows a DSNAOINI file with values based on the examples used in this section. Items in the file that may need to be customized to your DB2 installation are in bold type. See your DB2 administrator for the values of these items for your installation.

```
;This is a comment line...
; Example COMMON stanza
[COMMON]
MVSDEFAULTSSID=DSN5
; Example SUBSYSTEM stanza for your DB2 subsystem name
[DSN5]
MVSATTACHTYPE=CAF
;MVSATTACHTYPE=RRSAF
PLANNAME=DSNACLI
; Example DATA SOURCE stanza for your data source
[LOC1]
AUTOCOMMIT=0
CONNECTTYPE=1
```

Figure 2. Sample DSNAOINI file

Choosing the MVSATTACHTYPE

The LDAP server can be set up to use either the Call Attachment Facility (CAF) or the Recoverable Resource Manager Services Attachment Facility (RRSAF) to access DB2. See *DB2 for OS/390 and z/OS: ODBC Guide and Reference* for more information about these choices.

Installing RACF for SDBM and native authentication

In order for your LDAP server to have access to RACF data, you must have RACF installed on your system and have a license for the z/OS Security Server. RACF is part of the z/OS Security Server. Refer to the following books for information on installing and configuring RACF:

- z/OS: Program Directory
- z/OS: RACF Security Administrator's Guide
- z/OS: Security Server RACF Migration

The RACF Subsystem function of RACF must be defined and activated to allow the LDAP server to communicate with RACF through the SDBM backend. See *z/OS*: Security Server RACF System Programmer's Guide for information.

Installing and setting up Policy Director and SAF for z/OS Policy Director support

In order for your LDAP server to provide z/OS Policy Director support, you must install and set up Policy Director. See *Policy Director Authorization Service for z/OS and OS/390: Customization and Use* for instructions. Policy Director support also uses the System Authorization Facility (SAF) interface which is part of the z/OS environment and is always present on your z/OS system. *z/OS: Security Server RACF System Programmer's Guide* provides more information on SAF.

Installing System SSL

In order for your LDAP server to provide SSL/TLS support, you must install z/OS Cryptographic Services System SSL and use STEPLIB, LPALIB, or LINKLIST to make their libraries available. See "Setting up for SSL/TLS" on page 43 and z/OS: System Secure Sockets Layer Programming for more information regarding SSL/TLS. Also, see "Setting up for SSL/TLS" on page 43 for details on configuring and using SSL/TLS with your LDAP server.

Installing OCSF and ICSF for password encryption

The LDAP server uses the Open Cryptographic Services Facility (OCSF) to provide MD5 and SHA hashing of user passwords in the TDBM backend. The LDAP server uses both OCSF and the Integrated Cryptographic Service Facility (ICSF) to provide DES encryption and decryption of user passwords. The LDAP server does not require OCSF or ICSF to provide crypt() level encryption of user passwords. If you plan to encrypt passwords using MD5 hashing, SHA hashing, or DES encryption, you must install and configure the appropriate facility, or facilities, along with the LDAP server.

OCSF

In preparation for installing OCSF for password encryption, be sure to set LIBPATH in the /etc/ldap/slapd.envvars file. To install and configure OCSF, refer to the configuration information in z/OS: Open Cryptographic Services Facility Application Programming. This contains instructions on how to set up the necessary security authorizations using RACF to use OCSF. OCSF must be configured so that the user ID under which the LDAP server runs can use OCSF services. It also contains information on the Program Control necessary for OCSF. This documentation also contains instructions on how to run the installation scripts necessary to use OCSF.

Note: Although both OCSF and ICSF come with the base feature of z/OS, in the United States and Canada, an additional OCSF Security Level 3 feature must be ordered. There is no charge for this feature.

ICSF

To install, configure, and activate ICSF, your processor must have hardware cryptographic support. All new processors have hardware cryptographic support, while some older processors optionally provided this support.

Two other services of ICSF needed for DES encryption in the LDAP server are the Key Generator Utility Program (KGUP) and the Cryptographic Key Data Set (CKDS). These are needed to generate and store the key and key label needed for DES encryption of user passwords. Refer to the information about managing cryptographic keys and using the Key Generator Utility Program in z/OS: ICSF Administrator's Guide for instructions on how to generate and store into CKDS a single-length data-encrypting key (also referred to as data key) for DES encryption and how to set up the necessary security authorizations when using RACF to protect use of the key. It is important to remember to refresh both CKDS and RACF after you make the changes. ICSF must be configured so that the user ID under which the LDAP server runs can use ICSF services.

Other parts of the ICSF book may be useful for general background information about ICSF and Cryptographic Keys.

Installing Kerberos

In order for your LDAP server to provide Kerberos support, you must install the Security Server Network Authentication and Privacy Service for z/OS which is the IBM implementation of Kerberos Version 5. See z/OS: Security Server Network Authentication Service Administration for more information regarding Kerberos.

A sample Kerberos configuration file is provided in **/etc/skrb**. Refer to **z/OS**: Security Server Network Authentication Service Administration for details on setting up this file.

Chapter 4. Configuring an LDAP server using the Idapcnf utility

The LDAP configuration utility, **Idapcnf**, simplifies and automates the LDAP server configuration process for TDBM and, optionally, SDBM. The following table shows where to find specific information about the LDAP configuration utility in this chapter.

Description	Page
Overview of how the LDAP configuration utility works and information for determining if the utility is appropriate for your LDAP server configuration	17
Details describing how to use the Idapcnf utility	20
Step-by-step instructions describing how to configure an LDAP server	24
Configuration confirmation	27
Advanced configuration options	27
LDAP server configuration for other z/OS components or products, such as Managed System Infrastructure (msys)	29

Overview of the LDAP configuration utility

The LDAP configuration utility helps you configure new LDAP server instances with minimal user interaction.

The LDAP configuration utility takes a profile file as input and generates a set of output members in a data set to facilitate an LDAP server configuration. The profile file is targeted for the System Administrator (or System Programmer) and the LDAP Administrator and it contains statements that must be updated with appropriate values. The LDAP configuration utility generates a series of JCL members, configuration files, and a procedure to start the LDAP server. The JCL jobs are segregated based on typical administrative roles in a z/OS installation and contain the required commands to configure the z/OS components used by the LDAP server. Each administrator is responsible for reviewing and submitting their JCL job. After all JCL jobs are submitted, each administrator is responsible for reviewing their job's output and addressing any errors that may have occurred. Once all JCL jobs have completed successfully, the LDAP server can be started.

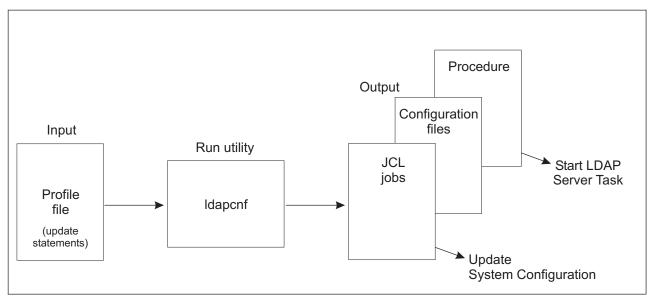


Figure 3. LDAP configuration utility overview

The minimal user interaction with the utility and the jobs it produces to update the required z/OS components results in a simplified approach to LDAP configuration. This approach allows novice LDAP users and administrators and even novice z/OS users to quickly deploy an LDAP server. In addition, the utility does not restrict the configuration of advanced LDAP features, such as referrals, replication. password encryption, and sysplex setup. See "Specifying advanced configuration options with the Idapcnf utility" on page 27 for more information.

Capabilities

Following are the capabilities of the LDAP configuration utility:

- Always configures a TDBM (DB2-based) backend and optionally configures an SDBM (RACF-based) backend. Also, Idapcnf can optionally configure an Extended operations (EXOP) backend.
- Generates JCL jobs to accomplish the updates of all the z/OS components required for an LDAP server.
- Can configure advanced LDAP server features, including:
 - Password encryption (Idapcnf does not generate certificates or passwords)
 - Referrals
 - Replication
 - Secure Sockets Layer (SSL) or Transport Layer Security (TLS) (Idaponf does not generate certificates or passwords)
 - Kerberos authentication
 - Native authentication
 - Extended operations (EXOP) backend (used for accessing Policy Directory information)

Restrictions

Following are some restrictions regarding the LDAP configuration utility:

- Generates a procedure; therefore, the LDAP server must run as a started task.
- Assumes that RACF is the security server in use. However, if RACF is not the security server in use, Idapcnf could still be used. The resulting RACF JCL job will need to be converted to properly update the security server in use.
- · Does not handle multiple TDBM (DB2-based) backends.
- · All values in the input files must be less than 66 bytes in length and must contain only printable characters in the IBM-1047 code page.
- · Cannot extend or enhance an existing LDAP server configuration. Furthermore, any manual updates to the output that the utility produces will be lost if you run the utility again with the same output data set.

- Does not support configuration for an LDAP server to listen on more than one secure port.
- Does not support configuration for an LDAP server to listen on more than one non-secure port.

If you cannot use the Idapcnf utility because of one or more of these restrictions, see Chapter 5, "Configuring an LDAP server without the Idapcnf utility" on page 31 for information on alternate methods you can use to configure your LDAP server.

Using the Idapcnf utility

The Idapcnf utility is used to generate jobs to set up the system environment and configuration for a new LDAP server. This utility is installed into the /usr/lpp/ldap/sbin directory.

Format

```
ldapcnf -i profile file
where:
profile file
        Specifies the input file that contains statements necessary to configure the LDAP server. See
        "Input file description" for more details about this file.
```

Example

Following is an example using **Idapcnf**, where **Idap.profile** is in the **/home/u** directory.

ldapcnf -i /home/u/ldap.profile

Input file description

The input file, Idap.profile, shipped in the /usr/lpp/ldap/etc directory, contains the settings necessary to set up an LDAP server. You must copy the Idap.profile file and then modify it before the LDAP configuration utility, Idapcnf, can be run.

In this file there are statements containing a keyword and value which must have the appropriate value for the target system being configured. Figure 4 shows a sample portion of the **Idap.profile** file:

```
# LDAPUSRID <user_id>
# Description:
    User ID for the LDAP server to run under.
    This variable's value must be capitalized.
LDAPUSRID="GLDSRV"
```

Figure 4. Sample portion of Idap.profile

The LDAPUSRID statement, as shown above, has a pre-assigned value of GLDSRV. Above the statement there is some commentary describing the statement and its usage.

Most of the statements in the Idap.profile are required and those that are not required are labelled as optional. Some statements in the Idap.profile have pre-assigned values; however, they may not be valid on the target system being configured. Values must be provided for all required statements in the Idap.profile file.

The Idap.profile file embeds three other advanced input files. Information about these files can be found in "Specifying advanced configuration options with the Idapcnf utility" on page 27. All of the input files are in the same format as an environment variable file.

Usage

1. The output from Idapcnf is written to an output data set that you specify in Idap.profile. If the data set does not exist, the utility allocates the output data set for you.

- 2. The utility allows the configuration of an LDAP server which uses SSL/TLS. (See "Setting up for SSL/TLS" on page 43 for details.) It does not, however, automate the process of generating SSL/TLS certificates.
- 3. The utility allows the configuration of an LDAP server which uses password encryption. It does not, however, automate the process of generating encryption keys. (See "Configuring for user password encryption" on page 47 for more information.)
- 4. It is recommended that the GLDHLQ.SGLDLNK data set containing the LDAP code be placed in LINKLST. If this is not done, then STEPLIB must be used to locate this data set.

The **Idapcnf** utility does not provide an interface for adding STEPLIB statements to the started task procedure that it generates. Therefore, if an administrator wishes to add STEPLIB statements to the started task procedure, the started task procedure must be manually updated and the following must occur:

- · The data sets specified in the new STEPLIB statements must be APF authorized.
- When submitting the PRGMCTRL JCL job, the data sets specified in the new STEPLIB statements must be in the program control data set list.
- The user ID specified on the LDAPUSRID statement in the Idap.profile file must have read access to the data sets specified in the new STEPLIB statements.
- 5. The APF JCL job will not work on a system using JES3. JES3 users need to manually enter the following operator command in place of submitting the APF JCL job:
 - In SDSF, enter:

/SET PROG=PROGsuffix

• From the operator's console, enter:

SET PROG=PROGsuffix

The *suffix* above is specified on the PROG SUFFIX statement in the **Idap.profile** file.

- 6. The PRGMCTRL and RACF jobs that **Idapcnf** generates require that the definitions listed below exist in RACF prior to submission. If the definitions do not exist, the jobs will contain RACF errors in their output.
 - a. To ensure that all required data sets are program controlled, the PRGMCTRL job requires that the PROGRAM ** definition exists in RACF.
 - b. To ensure that the user ID that appears on the LDAPUSRID statement in the Idap.profile file has read permission on all required data sets, the RACF job requires that data set definitions exist for the following data sets in RACF.
 - CEEHLQ.** (where CEEHLQ appears on the CEEHLQ statement in the Idap.profile file)
 - GLDHLQ.** (where GLDHLQ appears on the GLDHLQ statement in the Idap.profile file)
 - GSKHLQ.** (where GSKHLQ appears on the GSKHLQ statement in the Idap.profile file)
 - DSNHLQ.** (where DSNHLQ appears on the DSNHLQ statement in the Idap.profile file)
 - CBCHLQ.** (where CBCHLQ appears on the CBCHLQ statement in the Idap.profile file)
 - OUTPUT_DATASET_HLQ.** (where OUTPUT_DATASET_HLQ is the first level high-level qualifier of the data set name that appears on the OUTPUT_DATASET statement in the Idap.profile file)

Note that **Idapcnf** does not parse multi-level high-level qualifiers for the primary high-level qualifier. Thus, Idapcnf requires that data set definitions in RACF use the full high-level qualifiers as specified on the high-level qualifier statements in the Idap.profile file.

Also note that the server will operate properly even if the definitions required by the RACF JCL job do not exist, given that the user ID that appears on the LDAPUSRID statement in the Idap.profile file has read permission on all required data sets.

- 7. Administrators with the appropriate authorizations must submit the JCL jobs generated by Idaponf on the target system.
- 8. When running Idapcnf from an rlogin session, if OUTPUT_DATASET has not been pre-allocated, the script will try to free it after allocation. The free may exit with a return code 12, indicating that it is not

- currently allocated. The rlogin environment implicitly frees the data set when it is allocated, which is the cause of this error. This error can be ignored if running under this environment.
- 9. If an error occurs when submitting an Idapcnf output JCL job or, if prior to submission, an administrator considers a value within the JCL job unsatisfactory, the administrator should not modify the JCL job directly. Instead, the administrator should update the appropriate profile files and perform all of the steps outlined in "Steps for configuring an LDAP server" on page 24 again.
 - To help determine the statements within a profile file that the administrator may need to update, at the top of every file generated by Idapcnf there is a listing of statements that Idapcnf used to create the output file. The administrator can use this listing to determine the exact statement within a profile file that should be updated. Note that when resubmitting all the JCL jobs Idapcnf creates, many times JCL jobs for other components may have errors due to the duplication of a previous update. These messages may be ignored.
- 10. If a statement's value requires a length greater than 65 within the generated SLAPDCNF, the LDAP Administrator can move the SLAPDCNF member out of the output data set into a data set where the record length is greater than 80 bytes and update the member in the new data set. Then, the System Administrator must update the CONFIG DD card in the generated procedure to point to the new data
- 11. If the utility is configuring an SDBM backend and password encryption, the PRGMCTRL JCL job will use a shell script in Itmp called gldOcsfApf.sh. pid (where pid is the process ID of which the Idapcnf shell was run under) which is generated and placed in /tmp by the Idapcnf utility. This shell script must be available on the target system where the PRGMCTRL JCL job is submitted. This file should not be deleted until the PRGMCTRL member is submitted.
 - Also note that under the same conditions, the generated PRGMCTRL JCL job causes two HFS output files to be created: /tmp/gldCmd1.out. pid and /tmp/gldCmd1.err. These files report information about the success of the gldOcsfApf.sh shell.
 - If there were no problems encountered in the shell script when the PRGMCTRL JCL job is submitted, the shell script and the two output files will be deleted.
- 12. The profile files do not replace the LDAP server configuration file; they are used to create an LDAP server configuration file to run the LDAP server.
- 13. Error messages result if required statements are not assigned values in the input files or if simplified syntax checking fails.
- 14. It is recommended that you make all updates through the input files, running the utility again to recreate the jobs. Otherwise, if the generated JCL jobs are manually updated, those updates will be lost if the utility is run again using the same output data set.
- 15. Be sure to use a different output data set than is currently being used by other LDAP servers.

Configuration roles and responsibilities

The output from the LDAP configuration utility consists of jobs and configuration files that finalize the LDAP server configuration. These jobs segregate z/OS updates based on typical administrative roles, allowing each administrator to control their component's updates. The typical administrative roles that are assumed to exist to configure an LDAP server are:

- System Administrator (or System Programmer)
- Database Administrator
- LDAP Administrator
- Security Administrator

Each administrator is responsible for updating input files in addition to reviewing and submitting jobs in the output members that the LDAP configuration utility produces for their component, as shown in Table 2 on page 23.

Table 2. LDAP configuration utility roles and responsibilities

Role	Responsibility	Input file name/type	Output members
System Administrator (or System Programmer)	APF authorization	Idap.profile (main)	APF PROGsuffix (suffix is specified on the PROG_SUFFIX statement in Idap.profile)
Database Administrator	DB2, CLI	Idap.db2.profile (advanced)	DBSPUFI DBCLI DSNAOINI
LDAP Administrator	LDAP server, Kerberos authentication, native authentication	Idap.slapd.profile (advanced)	user_id procedure (user ID is specified on the LDAPUSRID statement in Idap.profile)
			SLAPDENV SLAPDONE
O it - A -l i i t t	DAOE 001/TLO	Lilan mark markita	• SLAPDCNF
Security Administrator	RACF SSL/TLS, password encryption	Idap.racf.profile Idap.slapd.profile (both files are advanced)	• RACF • PRGMCTRL

Note about Security Administrator: If configuring SDBM and password encryption, the Security Administrator must have read/write authority on all files in the /usr/lpp/ocsf/lib and /usr/lpp/ocsf/addins directories.

Figure 5 on page 24 is a graphical representation showing the administrative roles, input files, and output members for **Idapcnf**.

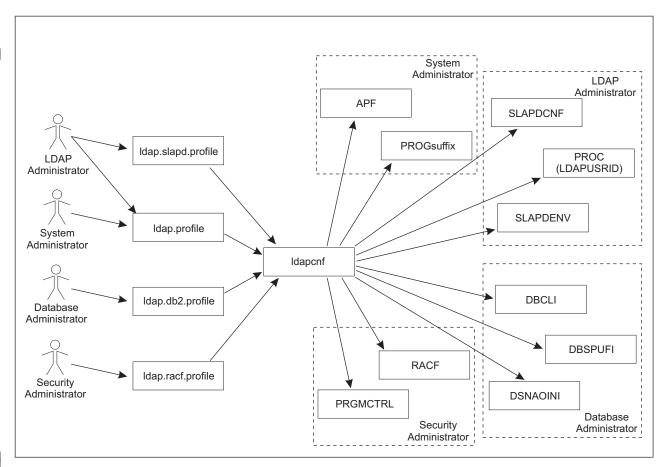


Figure 5. LDAP configuration utility roles and responsibilities

Steps for configuring an LDAP server

Use the following steps to configure with the configuration utility:

1. Copy the Idap.profile file, found in /usr/lpp/Idap/etc, to a local directory and update it according to the commentary found in the file. (If you need to update the advanced configuration files mentioned in Table 2 on page 23, you need to copy those files as well. See "Specifying advanced configuration options with the Idapcnf utility" on page 27 for more details on these files.)

Some statements in Idap.profile do not have any pre-assigned values but are required for successful configuration. These are noted in the file. Assign values to all of these required statements, referring to information in the file above each statement for assistance. The intended audience of the Idap.profile file is the System Administrator (or System Programmer) and the LDAP Administrator. The file contains information required from both administrators.

Note: Some statement values are case-sensitive and are denoted accordingly. Be sure to set up the editor to allow both upper and lower case letters to be specified.

2. Run the Idapcnf utility to generate members. (See "Using the Idapcnf utility" on page 20 for detailed information.) The generated members will be placed in the output data set specified on the OUTPUT DATASET statement in the **Idap.profile** file.

The utility creates:

- · JCL jobs for each role
- SLAPDCNF member which is the LDAP server configuration file
- SLAPDENV member which is the LDAP server environment variable file
- PROG member needed for APF authorization

- Procedure needed to start the LDAP server.
- · DSNAOINI configuration file for DB2 CLI
- DBSPUFI DB2 SQL statements
- 3. Copy members and submit jobs.
 - a. Copy the LDAP server started task procedure from the output data set to the target system's procedure library. The name of the LDAP server started task procedure will be the name of the LDAP user ID specified on the LDAPUSRID statement in the **Idap.profile** file. The pre-assigned name of the LDAP user ID is **GLDSRV**.
 - b. Copy the generated PROG*suffix* member (where *suffix* is specified on the PROG_SUFFIX statement in the **Idap.profile** file) from the output data set to the target system's PARMLIB.
 - c. Submit the following generated JCL jobs that can be found in the output data set in the following order:
 - 1) RACF member
 - 2) APF member
 - 3) DBCLI member

Note: Be sure DB2 is started before submitting this job.

4) PRGMCTRL member

The PRGMCTRL member is only required if one of the following scenarios exist:

- · Password encryption is being configured or
- · An SDBM backend is being configured and
 - Program Control is active
- 4. Through the DB2 SPUFI interactive tool, submit the DBSPUFI member.
- 5. Start the LDAP server. The LDAP server can be started from SDSF or from the operator's console.

Note: The name of the LDAP server procedure is the same as the user ID specified on the LDAPUSRID statement. The pre-assigned value is **GLDSRV**.

To start the LDAP server in SDSF, enter:

```
/s user_id
```

To start the LDAP server from the operator's console, enter:

```
s user id
```

- 6. Finalize set-up of the LDAP server.
 - a. Modify the schema entry for the TDBM backend.

Copy the **schema.user.ldif** file, found in the **/usr/lpp/ldap/etc** directory, into a local directory. Fill in the TDBM suffix in the line (shown below) which is near the top of the **schema.user.ldif** file.

```
dn: cn=schema, <suffix>
```

The suffix is the same value used in the TDBM_SUFFIX statement in the **Idap.profile** file. Here is an example with the suffix filled in:

```
dn: cn=schema, o=Your Company
```

Use the following **Idapmodify** utility (z/OS shell version) and the parameters shown to modify the schema entry.

```
\textbf{1dapmodify -h} \ \textit{ldaphost -p} \ \textit{ldapport -D} \ \textit{binddn -w} \ \textit{passwd -f} \ \textit{file}
```

where:

Idaphost

Is the host name of the system the LDAP server is running.

Idapport

Is the TCP/IP port on which the LDAP server is running. The port is specified in an advanced profile file, Idap.slapd.profile, on the PORT statement. The pre-assigned value is 3389.

binddn Is the administrator DN of the LDAP server. The administrator DN was specified in the Idap.profile file on the ADMINDN statement. The pre-assigned value is "cn=LDAP Administrator".

passwd

Is the administrator password of the LDAP server. The administrator password was specified in the Idap.profile file on the ADMINPW statement. The pre-assigned value is "secret".

file Is a file containing modifications to the schema entry in LDIF format. More information about the schema can be found in Chapter 16, "LDAP directory schema for TDBM" on page 159.

Following is an example of using **Idapmodify** to modify the schema entry:

ldapmodify -h myhost -p 3389 -D "cn=LDAP Administrator" -w secret -f /home/u/schema.user.ldif

This example assumes the schema file was copied to the /home/u directory and updated.

More information about Idapmodify can be found in z/OS: Security Server LDAP Client Programming.

Multiple schemas may need to be loaded before applications that use the directory will work. For example, in addition to the schema.user.ldif schema file, it is common for directory applications to require the elements defined in the schema.IBM.Idif schema file.

b. Load the suffix entry for the TDBM backend. (The suffix entry is specified in the Idap.profile file on the TDBM_SUFFIX statement.)

Notes:

- 1) This step can be ignored if, once the LDAP server has been started, the LDAP Administrator intends to load data into the directory from an LDAP Data Interchange Format (LDIF) file that contains the suffix entry.
- 2) If you intend to load large amounts of data in LDIF format into the directory, see "Idif2tdbm" program" on page 127 for instructions on using the **Idif2tdbm** utility. In this case, do not load the suffix entry separately. Include the suffix instead with the rest of the entries to be loaded by ldif2tdbm.

Use the following **Idapadd** utility (z/OS shell version) and the parameters shown to load the suffix

1dapadd -h ldaphost -p ldapport -D binddn -w passwd -f file

where:

Idaphost

Is the host name of the system the LDAP server is running.

Idapport

Is the TCP/IP port on which the LDAP server is running. The port is specified in an advanced profile file, Idap.slapd.profile, on the PORT statement. The pre-assigned value is 3389.

binddn Is the administrator DN of the LDAP server. The administrator DN is specified in the Idap.profile file on the ADMINDN statement. The pre-assigned value is "cn=LDAP Administrator".

passwd

Is the administrator password of the LDAP server. The administrator password was specified in the **Idap.profile** file on the ADMINPW statement. The pre-assigned value is "secret".

file Is a file containing the suffix entry in LDIF format. The distinguished name of the suffix entry must equal the value that appears on the TDBM_SUFFIX statement in the **Idap.profile** file. More information about LDIF format can be found in "Using LDIF format to represent LDAP entries" on page 290.

Following is an example of using **Idapadd** to load the suffix entry:

```
ldapadd -h myhost -p 3389 -D "cn=LDAP Administrator" -w secret -f suffix.ldif
```

More information about **Idapadd** can be found in *z/OS: Security Server LDAP Client Programming*.

To confirm the LDAP server is configured and ready for client requests, see "Configuration confirmation".

To load the data in LDIF format into the directory, you can use **Idif2tdbm** or **Idapadd**. However, if you intend to load more than 100,000 directory entries, use **Idif2tdbm**.

Configuration confirmation

Following is an optional Installation Verification Procedure (IVP) to confirm that the LDAP server configuration was successful.

Run the **Idapsearch** utility (z/OS shell version) with the following parameters to verify the configuration. **Idapsearch** -h *Idaphost* -p *Idapport* -V 3 -s base -b "" "objectclass=*"

where:

Idaphost

Is the host name of the system the LDAP server is running.

Idapport

Is the TCP/IP port on which the LDAP server is running. The port is specified in an advanced profile file, **Idap.slapd.profile**, on the PORT statement. The pre-assigned value is 3389.

The **-V 3** specifies LDAP Version 3 protocol, the **-s base** specifies the base scope for the search, and the **-b** "" specifies the root DSE as the base.

The result of this search is a list of all naming contexts supported by the LDAP server. For example, if both TDBM and SDBM are configured, the result of the search is both naming contexts (suffixes) listed.

Following is an example using **Idapsearch** to verify a configuration:

```
ldapsearch -h myhost -p 3389 -V 3 -s base -b "" "objectclass=*"
```

If the naming context is not returned, an error message is returned indicating a problem.

More information about Idapsearch can be found in z/OS: Security Server LDAP Client Programming.

Specifying advanced configuration options with the Idapcnf utility

There are advanced configuration options specified in the following input files:

- Idap.db2.profile (DB2 input file)
- Idap.racf.profile (RACF input file)
- Idap.slapd.profile (SLAPD input file)

These advanced profile files are located in the /usr/lpp/ldap/etc directory and are all included by the Idap.profile file. Every statement contains a pre-assigned value in the advanced profile files.

To modify these optional statements:

- 1. Copy the desired files to a local directory and update them. Each input file should be modified by the appropriate administrator (see Table 2 on page 23).
- 2. Update the Idap.profile to correctly include those modifications. Near the end of Idap.profile there are three statements:

```
${SOURCE CMD} ${USR LPP ROOT}/usr/lpp/ldap/etc/ldap.slapd.profile
${SOURCE CMD} ${USR LPP ROOT}/usr/lpp/ldap/etc/ldap.db2.profile
${SOURCE_CMD} ${USR_LPP_ROOT}/usr/lpp/ldap/etc/ldap.racf.profile
```

Update these statements to show the new paths of the files you modified. Here is an example where the modified versions of Idap.slapd.profile and Idap.db2.profile are in different directories (/home/u), and **Idap.racf.profile** was not changed:

```
${SOURCE CMD} /home/u/ldap.slapd.profile
${SOURCE_CMD} /home/u/ldap.db2.profile
${SOURCE CMD} ${USR LPP ROOT}/usr/lpp/ldap/etc/ldap.racf.profile
```

Advanced configuration options may require additional instructions not covered by the LDAP configuration utility. The following table provides references for those instructions.

Configuration option	More information	
Referrals	Chapter 25, "Referrals" on page 277	
Replication	Chapter 24, "Replication" on page 269	
Password encryption	"Configuring for user password encryption" on page 47	
Multi-server	"Determining operational mode" on page 75	
Kerberos authentication	Chapter 19, "Kerberos authentication" on page 225	
Native authentication	Chapter 20, "Native authentication" on page 233	
CRAM-MD5 and DIGEST-MD5 authentication	Chapter 21, "CRAM-MD5 and DIGEST-MD5 Authentication" on page 243	
Extended operations to access Policy Director data	"Setting up for extended operations" on page 43	
Entry UUID support	serverEtherAddress option on page 66	
Other LDAP server options	Chapter 8, "Customizing the LDAP server configuration" on page 51	

Notes:

- 1. If the UID in the Idap.racf.profile (specified on the on LDAPUID statement) is greater than 0 and the port or the secureport in the Idap.slapd.profile file (specified on the PORT and SECUREPORT statements, respectively) is less than 1024, the PORT statements generated in the SLAPDCNF member commentary must be added to the **profile.tcpip** file on the target system. These PORT statements are located in the commentary directly above the port and secureport variables in the generated SLAPDCNF member. See z/OS: Communications Server: IP Configuration Guide for more information on the profile.tcpip file.
- 2. When using Idapcnf to configure an LDAP server with Kerberos enabled, the user ID that the server runs under is created with a temporary password. This temporary password is then immediately removed. This is required to complete the configuration of an LDAP server with Kerberos enabled. See Chapter 19, "Kerberos authentication" on page 225 for more details.

LDAP server configuration for other z/OS components or products

The LDAP Configuration Utility can be used to configure an LDAP server for use by other z/OS components or products, such as Managed System Infrastructure (msys). Some assigned values in the input file, or the advanced input files, may not be valid for an LDAP server required for other z/OS components or products.

For msys, there is a set of input files available that can be used to configure an LDAP server which will be used by msys. The input files can be found in /usr/lpp/ldap/etc and have the following names:

- · Idap.msys.profile
- Idap.msys.db2.profile
- Idap.msys.racf.profile
- · Idap.msys.slapd.profile

The **Idap.msys.profile** file corresponds to the **Idap.profile** file described in the previous sections in this chapter. Substitute **Idap.msys.profile** for **Idap.profile** in the instructions in "Using the Idapcnf utility" on page 20 and "Steps for configuring an LDAP server" on page 24 to modify **Idap.msys.profile**.

In addition to loading the **schema.user.ldif** file (explained in step 6 on page 25), you must also load the **schema.IBM.ldif** for msys.

Chapter 5. Configuring an LDAP server without the Idapcnf utility

This chapter lists the necessary steps involved in configuring your LDAP server if you do not use the **Idapcnf** utility. It may be necessary for you to use this method instead of the **Idapcnf** utility, as all LDAP configuration scenarios cannot be set up with the **Idapcnf** utility. More information about the **Idapcnf** utility is in Chapter 4, "Configuring an LDAP server using the Idapcnf utility" on page 17.

This chapter contains:

- A roadmap which provides LDAP server configuration steps based on which backends and options you choose to configure, such as:
 - SDBM backend (RACF-based)
 - TDBM backends (DB2-based)
 - EXOP backend for accessing Policy Director data
 - Secure Sockets Layer (SSL) or Transport Layer Security (TLS)
 - Password encryption
 - Kerberos authentication
 - Native authentication
- · A list of configuration variables and their interactions

Using TDBM provides the following advantages:

- Initial load times extremely faster with Idif2tdbm
- All database definitions in SPUFI
- · Simple, fully-indexed search table
- Supports large numbers of entries
 - · Has a flexible method of defining overflow data
 - · Stores data in UTF-8 format, avoiding translation overhead
 - · Robust schema support, including:
 - Schema publication
 - Dynamic schema update support
 - Compliance with matching rule definitions
 - Many syntaxes supported
 - · Support of multiple TDBM backends
 - · Provides exception reporting, including detail on diagnostic errors returned to the client
 - · Provides enhanced ACL support:
 - cn=authenticated (to allow non-specific access beyond what cn=anybody would allow)
 - cn=this (to allow individuals access to their own entry, with the ACL defined once across an entire sub-tree)
 - Ability to add or remove individual ACL entries with Idapmodify

LDAP server configuration roadmap

Table 3 lists the set up and configuration tasks you must complete depending on which backends and options your LDAP server needs.

Table 3. LDAP server configuration roadmap

Complete?	Task	Page	
If you are confi	If you are configuring an SDBM (RACF-based) backend, you must:		
	Install RACF	15	
	Set up the User ID and Security for the LDAP server	35	
	Set up the LDAP server for SDBM	42	
	Configure the LDAP server	51	

1

Table 3. LDAP server configuration roadmap (continued)

Complete?	Task	Page			
	Run the LDAP server	93			
	See other SDBM-specific information.	205			
If you are conf	figuring a TDBM (DB2-based) backend, you must:	-			
	Install and set up DB2	13			
	Set up the User ID and Security for the LDAP server.	35			
	Create the DB2 database and table spaces for TDBM	40			
	Set up the schema for TDBM	159			
	Load the data into the LDAP server for TDBM	127			
	Configure the LDAP server	51			
	Run the LDAP server	93			
If you are con	nfiguring an EXOP (extended operation) backend, you must:				
	Install Policy Director	15			
	Set up the User ID and Security for the LDAP server	35			
	Set up the LDAP server for EXOP	43			
	Configure the LDAP server	51			
	Run the LDAP server	93			
If your LDAP s	server is going to support Secure Sockets Layer (SSL) or Transport Layer Secur	rity (TLS), you must:			
	Install and set up System SSL	15			
	Locate System SSL and protect the environment for use of SSL/TLS	37			
	Set up the LDAP server for SSL/TLS	43			
	Configure the LDAP server	51			
	Run the LDAP server	93			
If your LDAP s	server is going to use password encryption, you must:				
	Install OCSF and ICSF	16			
	Protect the environment for OCSF and ICSF	37			
	Set up the LDAP server for password encryption	47			
	Configure the LDAP server	51			
	Run the LDAP server	93			
If your LDAP s	server is going to support Kerberos Authentication, you must:				
	Install and configure Kerberos	16			
	Start the KDC	225			
	Update the Kerberos segment of the LDAP server's user ID	225			
	Generate the LDAP server's key table file (optional)	225			
	Configure the LDAP server	51			
	Run the LDAP server	93			
If your LDAP s	server is going to support native authentication, you must:	'			
	Install RACF or other security server	15			
	Configure the LDAP server	51			
	Run the LDAP server	93			

Preparing for configuration variable interactions

Some of the variables involved in configuring the LDAP server and its related products are used in more than one file or configuration step. It is essential that the same value be used each time the variable is referenced. The following table lists the interactions of each such variable, for each backend.

Table 4. Configuration variable interactions

Variable	Used in:
TDBM Backend	
DB2 subsystem ID	SYSTEM(DSN) value in CLI bind JCL, DSNTIJCL (see Step 3 on page 14) MVSDEFAULTSSID and SUBSYSTEM values in CLI initialization file, DSNAOINI (see Step 4 on page 14)
Plan name	 PLAN value in CLI bind JCL, DSNTIJCL (see Step 3 on page 14) PLANNAME value in CLI initialization file, DSNAOINI (see Step 4 on page 14) The zzz value in SQL commands to grant permissions to LDAP user (see Step 5 on page 41)
Database name	 -DDDDDDDD- value in SPUFI script to create database, <i>GLDHLQ</i>.SGLDSAMP(TDBMDB) (see Step 2 on page 40) The <i>databasename</i> value in LDAP configuration file, slapd.conf (see Page 56) The <i>yyy</i> value in SQL commands to grant permissions to LDAP user (see Step 5 on page 41)
Database owner	 -UUUUUUUU- value in SPUFI script to create database, <i>GLDHLQ</i>.SGLDSAMP(TDBMDB) (see Step 2 on page 40) The <i>dbuserid</i> value in LDAP configuration file, slapd.conf (see Page 56)
CLI initialization file name	 The dsnaoini value in LDAP configuration file, slapd.conf (see Page 56) DSNAOINI DD value in JCL for LDAP server and utilities (see Page 413)
User ID running LDAP server, tdbm2ldif, or ldif2tdbm	 LDAPSRV value in RACF commands to create user ID (see Page 35) The xxx value in SQL commands to grant permissions to LDAP user (see Step 5 on page 41) LDAPSRV value in RACF commands to create SSL/TLS key ring (see Page 45) LDAPSRV value in RACF commands to create started task (see Page 93)
Server name	DATA SOURCE value in CLI initialization file, DSNAOINI (see Step 4 on page 14) The servername value in LDAP configuration file, slapd.conf (see Page 67)
SDBM Backend	
User ID running LDAP server or utilities	LDAPSRV value in RACF commands to create user ID (see Page 35) LDAPSRV value in RACF commands to create SSL/TLS key ring (see Page 45) LDAPSRV value in RACF commands to create started task (see Page 93)

33

Chapter 6. Setting up the user ID and security for the LDAP server

This chapter discusses how to configure and set up the products needed by your LDAP server.

If you plan to use:	You must:	See:
LDAP server as a started task	Define the LDAP server user ID.	"Setting up a user ID for your LDAP server" below and continue through remaining sections of this chapter.
LDAP server from USS shell	Set up user ID and environment.	"Requirements for a user ID that runs the LDAP server" and continue through remaining sections of this chapter.

Setting up a user ID for your LDAP server

When running the LDAP server as a started task, it is recommended that a separate user ID be established for the LDAP server. This section describes how to define the user ID that runs the LDAP server.

In this section, some of the examples and descriptions reflect assumptions that may not apply to your environment. Following are descriptions of these assumptions, with guidance on how to use this information if they do not apply to your environment:

- Some examples use Resource Access Control Facility (RACF). You can use any z/OS external security
 manager that has equivalent support. You must substitute the appropriate procedures for any examples
 that use RACF.
- The default name /usr/lpp/ldap is used for the directory in which you installed the LDAP server product. If you used a different name, substitute that name in the examples and descriptions where applicable.
- The language setting En_US.IBM-1047 is used for the locale in which you are running the LDAP server. This setting is used in the names of several directories that are referred to in this information. If you are using a different language setting, substitute that setting in the examples and descriptions where applicable. You must also specify this setting as the value of the LANG parameter in the environment variable file as described in Chapter 14, "Internationalization support" on page 151. The default environment variable file already sets LANG to En_US.IBM-1047.
- The name **LDAPSRV** is used for the user ID that runs the LDAP server. If you use a different name, substitute that name in the examples and descriptions where applicable.
- The name of the production directory is /etc/ldap. If you use a different name, you must symbolic link the names of the appropriate files in your directory to the /etc/ldap directory.

Requirements for a user ID that runs the LDAP server

Any user ID can be used to run the LDAP server. The examples in this chapter use a user ID of LDAPSRV in the commands provided.

The user ID that runs the LDAP server must have the following attributes:

- If you defined the BPX.DAEMON facility class, the user ID must have read access to the class.
- If you defined the BPX.SERVER facility class, the user ID must have update access to the class.
- The user ID must have read access to the data sets defined in the startup procedure

Note that if the UID of the user ID running the LDAP server is not zero, all console messages produced by the LDAP Server are accompanied by a BPXM023I message identifying the user writing to the console.

	The user ID performing the RACF commands in the following examples requires RACF SPECIAL authority.
	You can use the RACF commands in the following example for the user ID that will run the LDAP server. ADDGROUP LDAPGRP SUPGROUP(SYS1) OMVS(GID(2)) ADDUSER LDAPSRV DFLTGRP(LDAPGRP) OMVS(UID(1) PROGRAM ('/bin/sh'))
	The following commands should only be entered if these facility classes are defined. PERMIT BPX.DAEMON CLASS(FACILITY) ID(LDAPSRV) ACCESS(READ) PERMIT BPX.SERVER CLASS(FACILITY) ID(LDAPSRV) ACCESS(UPDATE)
	If you are going to set up more than one LDAP server on the same system, a separate user ID should be used for each one.
	Additional setup when using SDBM If you plan to use an SDBM backend, the following RACF commands must be entered to set up the user ID that will run the LDAP server: RDEFINE FACILITY IRR.RUSERMAP UACC(NONE) PERMIT IRR.RUSERMAP CLASS(FACILITY) ID(LDAPSRV) ACCESS(READ) SETROPTS RACLIST(FACILITY) REFRESH
	The SDBM backend also supports the RACF functions that search for users and groups with a given UID or GID value and control sharing user UID and group GID values. Usage of these functions requires additional RACF configuration and profiles, as described in the RACF documentation.
	Additional setup for RACF PROXY segment and SDBM The SDBM backend supports the PROXY segment within the RACF user profile. If you intend to use SDBM to set the BINDPW value in the PROXY segment, RACF requires that the KEYSMSTR class profile LDAP.BINDPW.KEY be created with the SSIGNON segment. • To create the LDAP.BINDPW.KEY profile in the KEYSMSTR class, use the KEYMASKED sub-operand if no cryptographic product is installed on your system: RDEFINE KEYSMSTR LDAP.BINDPW.KEY SSIGNON(KEYMASKED(key-value))
	Or, use the KEYENCRYPTED sub-operand if a cryptographic product is installed: RDEFINE KEYSMSTR LDAP.BINDPW.KEY SSIGNON(KEYENCRYPTED(key-value))
	key-value is a Secured Sign-on Application key and must be specified as a string of 16 hexidecimal characters.Then, activate the KEYSMSTR class: SETROPTS CLASSACT (KEYSMSTR)
	See <i>z/OS:</i> Security Server RACF Command Language Reference for details on using these RACF commands and <i>z/OS:</i> RACF Security Administrator's Guide for information on creating and using profiles.
	Defining the Kerberos identity If you plan to enable Kerberos support you need to associate a Kerberos identity with the server's user ID. The following command must be entered: ALTUSER LDAPSRV PASSWORD(password) NOEXPIRED KERB(KERBNAME(LDAP/hostname))
 	If the server is being run as a started task, also enter the following command: ALTUSER LDAPSRV NOPASSWORD
	Note that the word "LDAP" must be uppercase in the KERBNAME section of the ALTUSER command. Also, the hostname needs to be the primary hostname for the system in DNS.

If the LDAP server is located on the same machine as the Key Distribution Center then a keytab file is not necessary to start the LDAP server. However the user ID that starts the server must have at least read access to IRR.RUSERMAP in the FACILITY class. This can be done by issuing the following commands:

```
RDEFINE FACILITY IRR.RUSERMAP UACC(NONE)
PERMIT IRR.RUSERMAP CLASS(FACILITY) ID(LDAPSRV) ACCESS(READ)
SETROPTS RACLIST(FACILITY) REFRESH
```

Protecting the environment for the LDAP server

The PDS which contains the LDAP server, GLDHLQ.SGLDLNK, and the PDSs containing all the DLLs that the LDAP server loads must be APF-authorized to allow the LDAP server to make the necesary program control threading calls. For example:

```
SETPROG APF, ADD, DSN=GLDHLQ. SGLDLNK, VOL=volid
```

Additionally, if program control is active on your system, the PDS which contains the LDAP server and DLLs, the PDS that contains the C runtime libraries, SYS1.LINKLIB, and SYS1.CSSLIB must be program controlled. Also, the PDS containing the DB2 CLI (Call Level Interface), DB2HLQ.SDSNLOAD, must be APF-authorized and program controlled. Also, if a decision has been made to use password encryption support in the TDBM backend, there are additional requirements for APF-authorization and program control. See "Protecting the environment for password encryption" for those requirements.

Protecting the environment for SSL/TLS

If you are using SSL/TLS to secure your LDAP server, the STEPLIB or LINKLIST must include GSKHLQ.SGSKLOAD. Also, GSKHLQ.SGSKLOAD must be APF-authorized.

Protecting the environment for password encryption

If the LDAP server is configured to encrypt passwords, the OCSF libraries must also be APF-authorized. Since the OCSF DLLs are stored in HFS, the APF-authorized extended attribute must be turned on for the OCSF DLLs. The DLLs (.dll and .so files) in the /usr/lpp/ocsf/lib and /usr/lpp/ocsf/addins directories must have their APF-authorized extended attribute turned on by using the extattr +a command. First, use the following commands:

```
RDEFINE FACILITY BPX.FILEATTR.APF UACC(NONE)
PERMIT BPX.FILEATTR.APF CLASS(FACILITY) ID(userid) ACCESS(UPDATE)
SETROPTS RACLIST(FACILITY) REFRESH
```

where userid is the ID from which the extattr command will be run. That user ID can now run the extattr command.

Then use the following commands from an OMVS command prompt:

```
cd /usr/lpp/ocsf/lib
extattr +a *.dll
cd /usr/lpp/ocsf/addins
extattr +a *.so
```

In this example, the \$ (dollar sign) represents the OMVS command prompt and is not part of the command.

Refer to z/OS: UNIX System Services Planning for more details.

OCSF DLLs must also be program controlled. See the configuration information in z/OS: Open Cryptographic Services Facility Application Programming for more information.

Chapter 7. Preparing the backends, SSL/TLS, and password encryption

This chapter discusses what you need to do to prepare the backends, SSL/TLS, and password encryption.

If you plan to use:	You must:	See:
A backend, such as TDBM or SDBM	Copy the configuration files.	"Copying the configuration files"
The sample server to set up a DB2 database	Use the set of example files shipped with the code.	"Creating a DB2 database for the sample server" below
TDBM backend	Create the DB2 database and table spaces using SPUFI.	"Creating the DB2 database and table spaces for TDBM" on page 40
SDBM backend	Set up your configuration files.	"Setting up for SDBM" on page 42
EXOP backend	Set up your configuration files.	"Setting up for extended operations" on page 43
SSL/TLS	Enable SSL/TLS support.	"Setting up for SSL/TLS" on page 43
Password encryption	Determine the type of encryption and set up the configuration.	"Configuring for user password encryption" on page 47

Copying the configuration files

The configuration files need to be copied from the directory in which they are installed, /usr/lpp/ldap/etc, to the directory where they are used, /etc/ldap. Do not modify these files in the install directory because any service to the files will overwrite the modifications. Instead, modify them in /etc/ldap. The following commands copy the configuration files:

cp /usr/lpp/ldap/etc/slapd.* /etc/ldap/.

The **cp** command creates a copy of the **/usr/lpp/ldap/etc** files into the **/etc/ldap** directory. The individual files can now be modified using **oedit** or **vi**.

Creating a DB2 database for the sample server

There is a set of example files shipped in /usr/lpp/ldap/examples/sample_server that can be used to understand how to configure and run the LDAP server using a TDBM backend. The **README** provides step-by-step instructions for getting an LDAP server configured and started quickly. The following list shows the files shipped in that directory.

- **README** (Installation information for sample server)
- dsnaoini.db2ini (Sample CLI initialization file)
- dsntijcl.jcl (Sample job to bind CLI packages)
- sample.ldif (Sample directory entry for sample server)
- Idaptbl.jcl (Sample JCL to create database for TDBM)
- tdbm_sample_db.spufi (Sample SPUFI script to create tables and table spaces for TDBM)
- tdbm_sample_index.spufi (Sample SPUFI script to create indexes for TDBM)

© Copyright IBM Corp. 1998, 2002

Creating the DB2 database and table spaces for TDBM

When using TDBM, the LDAP server DB2 database must be created by running two SPUFI (SQL Processor Using File Input) scripts from DB2 Interactive (DB2I), DB2I is a DB2 facility that provides for the running of SQL statements, DB2 (operator) commands, and utility invocation. For details on how to use DB2I and SPUFI, see DB2 for OS/390 and z/OS: Application Programming and SQL Guide. Sample DB2I SPUFI scripts to create the LDAP server DB2 database are provided. To use them, do the following:

1. Copy the SPUFI scripts over to your SPUFI input data set.

The SPUFI script for creating the database and table spaces can be found in GLDHLQ.SGLDSAMP(TDBMDB) and the script for creating the table indexes can be found in GLDHLQ.SGLDSAMP(TDBMINDX). (GLDHLQ refers to the high-level qualifier that was used to install the LDAP server data sets.)

2. Determine values for SPUFI script.

In order to create the DB2 database and table spaces for TDBM, you must first decide on certain values within these SPUFI files, as shown in Table 5. (Table 4 on page 33 lists variables that are used in more than one file or configuration step. Be sure to specify the same values where necessary.) The SPUFI scripts provide specific instructions and information to help you determine the values to use in the table. "The TDBMDB SPUFI file" on page 403 and "The TDBMINDX SPUFI file" on page 409 show examples of the files to edit and run in the SPUFI facility.

Table 5. TDBM value overview

Attribute	Value	Suggested value	Variable name in SPUFI script	
Database information for TDBMDB and TDBMINDX members				
Database name		LDAPSRV	-DDDDDDDD-	
Database owner		LDAPSRV	-UUUUUUU-	
Table space definitions for TDBMDB me	mber			
Entry table space name		ENTRYTS	-AAAAAAA-	
Buffer pool name for the LDAP entry table space		BP0	-BBBB-	
Long entry table space name		LENTRYTS	-CCCCCCC-	
Buffer pool name for the LDAP long entry		BP0	-DDDD-	
Long attribute table space name		LATTRTS	-EEEEEEE-	
Buffer pool name for the LDAP long attribute		BP0	-FFFF-	
Miscellaneous table space name		MISCTS	-GGGGGGG-	
Search table space name		SEARCHTS	ННННННН-	
Buffer pool name for the LDAP search table		BP0	-1111-	
Replica table space name		REPTS	-JJJJJJJ-	
Descendants table space name		DESCTS	-KKKKKKKK-	
Storage group		SYSDEFLT	-SSSSSSS-	
Search column truncation size (VALUE in DIR_SEARCH)		32	-ТТТТ-	
DN truncation size (DN_TRUNC in DIR_ENTRY)		32	-MMMM-	
Maximum size of a DN (DN in DIR_ENTRY)		512	-NNN-	

3. Modify the scripts.

Use the values from Table 5 on page 40 to modify the scripts.

4. Run the scripts from DB2I SPUFI.

Use the DB2 SPUFI (SQL Processor Using File Input) facility to create the database and table spaces. Be sure to run the two scripts that were copied and modified in the previous steps under a user ID with DB2 **SYSADM** authority. When the scripts complete running, scan the output data set to ensure that they ran successfully.

5. Grant appropriate DB2 resource authorizations.

In order to run the LDAP server and server utilities, certain minimum DB2 resource authorizations must be granted to the user ID or user IDs that will be running these programs. Following are the suggested minimums which should be granted to those user IDs, where xxx is the user ID running the LDAP server or LDAP server utility, yyy is the database name identified in the **slapd.conf** and SPUFI file for the **databasename** option and zzz is the CLI plan name as specified in your DB2 CLI initialization file. Run the following statements through SPUFI (DB2 Interactive):

```
grant execute on plan zzz to xxx;
grant dbadm on database yyy to xxx;
```

These privileges may be granted by any user ID with **SYSADM** authority. The commands above can be run using the DB2 SPUFI facility.

The LDAP server requires **SELECT** access to the SYSIBM.SYSCOLUMNS table in DB2. If **SELECT** access to this table is tightly controlled in your DB2 installation, then it may be necessary to grant this access to the user ID under which the LDAP server runs by performing the following operation (either through SPUFI or another means of issuing SQL commands):

```
grant select on sysibm.syscolumns to xxx;
```

where xxx is the user ID under which the LDAP server runs. If this authority is not granted to the user ID under which the LDAP Server runs, the LDAP server will fail during start-up with an SQL -551 return code.

Partitioning DB2 tables for TDBM

If you are creating a large directory, you should partition the "entry table space" and "search table space" to improve performance and ease maintainability of the database. The following information identifies the partitioning indexes and values to use when partitioning these table spaces.

Table 6. TDBM table space partitioning indexes and values

Table space	Table name	Partitioning index	Partitioning column	Value range of column
Search tablespace	DIR_SEARCH	DIR_SEARCHX2	EID	0-9999999999999
Entry tablespace	DIR_ENTRY	DIR_ENTRYX0	EID	0-9999999999999

The EID value generated by the TDBM backend is a random 15 digit decimal number between 1 and 9999999999999. To determine the maximum value to assign to each partition, use the following formula:

```
eids_per_partition = 1000000000000000 / number_of_partitions
partition_value = eids_per_partition * partition_number
```

You can also partition the following additional table spaces in TDBM using the EID range as the partitioning value

Table 7. TDBM table space partitioning using EID range

Table space	Table name	Partitioning index	Partitioning column	Value range of column
Descendants table space	DIR_DESC	DIR_DESCX1	DEID	0-99999999999999
Long attribute table space	DIR_LONGATTR	DIR_LONGATTRX1	EID	0-99999999999999
Long entry table space	DIR_LONGENTRY	DIR_LONGENTRYX1	EID	0-9999999999999

Partitioning example

If your directory is to contain 10 million entries and you want 1 million entries per partition, the maximum value for each partition is:

Partition number	Partition maximum value
1	10000000000000
2	20000000000000
3	30000000000000
4	40000000000000
5	50000000000000
6	60000000000000
7	70000000000000
8	8000000000000
9	90000000000000
10	9999999999999

Setting up for SDBM

The LDAP server can provide LDAP access to the user and group information stored in RACF. See Chapter 18, "Accessing RACF information" on page 205 for details about how you can use this RACF information.

In order to configure your LDAP server to run with the SDBM database of the LDAP server:

- If you have not already done this, copy the configuration files from the /usr/lpp/ldap/etc directory to the /etc/Idap directory (see "Copying the configuration files" on page 39).
- · You need to use the following lines in your slapd.conf file:

```
database sdbm GLDBSDBM
suffix "your_suffix"
```

where your_suffix is any valid DN (distinguished name). Be sure to provide a meaningful value for the suffix. Note that it is no longer required that the **sysplex** attribute be present in the suffix. For example, a valid suffix line is:

```
suffix "cn=RACFA,o=IBM,c=US"
```

An SDBM suffix should not contain an alias name for an attribute. For example, the suffix cannot use the **surName** attribute (it can use the **sn** attribute instead). Also, the suffix can contain a case-sensitive attribute, but SDBM ignores case when processing the suffix.

 Be sure to set STEPLIB to point to the GLDHLQ.SGLDLNK data set before running the LDAP server with SDBM, otherwise an error will be returned. (GLDHLQ refers to the high-level qualifier that was used to install the LDAP server data sets.)

Notes:

- 1. Only one SDBM database can be defined in any given LDAP server.
- 2. SDBM contains an internal schema that it uses to check entries that it is adding. The schema cannot be modified.

Running SDBM with other backends

The following table gives you information on running SDBM alone or with other backend databases.

Table 8. SDBM with other backends

Backend	Description
SDBM with TDBM	The TDBM schema will be used for all initial DN normalization if TDBM is configured. DN normalization is performed by the server to aid in selecting the appropriate backend. All attribute types that might appear in a RACF-style DN must be defined to the TDBM schema. When starting TDBM and SDBM together, ensure that the attribute types in the SDBM suffix
	are also present in the TDBM schema. Examine the schema LDIF files shipped with the LDAP server to determine which schema must be loaded into TDBM.
SDBM only	If you are running SDBM without TDBM, be sure to comment out the TDBM database definitions in the slapd.conf file. Prefix each line to comment with a # (pound sign). When running only an SDBM backend, replication and referrals are not supported.

Setting up for extended operations

The LDAP server supports extended operations (EXOP backend) that retrieve Policy Director data. See Chapter 22, "Using extended operations to access Policy Director data" on page 247 for details on using this extended operations support.

To configure your LDAP server to run with the EXOP backend:

- Make sure your slapd.conf configuration file contains the following line for Policy Director: listen ldap://:pc
- Make sure your slapd.conf configuration file contains the following line: database exop GLDXPDIR

More information about the configuration file and options is in Chapter 8, "Customizing the LDAP server configuration" on page 51.

Setting up for SSL/TLS

The LDAP server contains the ability to protect LDAP access with Secure Sockets Layer (SSL) and Transport Layer (TLS) security. There are two types of connections that support secure communication:

- · An SSL/TLS only secure connection. This connection requires that the first communication between the client and the server be the handshake that negotiates the secure communication. From that point on only secure communication can occur on the connection.
- A bimodal connection that supports secure and non-secure communication. The client is expected to begin communication in a non-secure mode. At some time during communication, the client may change to secure communication by sending a StartTLS extended operation after which the handshake to negotiate secure communication occurs followed by secure communication. The client may shutdown secure communication causing a StopTLS alert to be sent and the server will continue communication in a non-secure mode. At a later time, the client may restart secure communication by sending another StartTLS extended operation followed by the handshake.

Both types of connections require that System SSL be configured for use by the LDAP Server.

Using SSL/TLS protected communications

- The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols use public-key infrastructure (PKI) algorithms to establish and maintain an encrypted communications path between a client and server. In z/OS, the ability to set up and communicate over SSL/TLS protected communication links is provided by the LDAP server with a set of services provided in z/OS (the z/OS Cryptographic Services System SSL set of services).
- In order for the LDAP client to communicate with an LDAP server over an SSL/TLS-protected TCP/IP socket connection, the LDAP server must transmit a certificate to the LDAP client and, optionally, the client can transmit its certificate to the LDAP server. The LDAP client and server must verify that the certificates they received are valid. Once the LDAP client and LDAP server have determined the validity of the certificates provided to them, SSL/TLS-protected communication occurs between the LDAP client and server.
- The LDAP client and server verify the certificates sent to them by using public-key digital signatures. The LDAP client and server take the certificates and compare the digital signature in the certificates with a signature that it computes based on having the public-key of the signer of the certificate. In order to do this, the LDAP client and server must have the public-key of the signer of the certificates. The LDAP client and server obtain this by reading a file that contains these public-keys. This file is called a key database.
- A key database, or RACF key ring, contains the public-keys that are associated with signers of certificates. These public-keys are, in reality, contained in certificates themselves. Thus, verifying one certificate requires the use of a different certificate, the signer's certificate. In this fashion, a chain of certificates is established, with one certificate being verified by using another certificate and that certificate being verified by yet another certificate, and so on. A certificate, and its associated public key, can be defined as a root certificate. A root certificate is self-signed, meaning that the public-key contained in the certificate is used to sign the certificate. Using a root certificate implies that the user trusts the root certificate.
- The key databases, or key rings, used by the LDAP client and server must contain enough certificates in order to verify the certificates sent by the LDAP client and server during the start-up of the SSL/TLS connection. If either certificate is self-signed, then that certificate must be stored in the other's key database. If the certificates are signed by some other certificate signer, then the signer's certificate and any certificates that this certificate depends upon must be stored in the key databases. The key databases used by the LDAP client and server must also contain the certificates that they will transmit to each other during the startup of the SSL/TLS-protected communications.

Enabling SSL/TLS support

The following high-level steps are required to enable SSL/TLS support for LDAP. These steps assume you have already installed and configured the LDAP directory server and installed z/OS Cryptographic Services System SSL and set STEPLIB, LPALIB, or LINKLIST.

- 1. Generate the LDAP server private key and server certificate and mark it as the default in the key database or use its label on the sslCertificate configuration file option (see "Using SSL/TLS protected communications").
- 2. Configure the LDAP server to listen for LDAP requests and configure the type of authentication wanted, server and optionally client authentication (see "Setting up the security options for the LDAP server" on page 45).
 - · For a secure only socket, a listen configuration or command line option must be set up for the secure port.
 - · For a bimodal socket, a listen configuration or command line option must be set up for the non-secure port.
- 3. Restart the LDAP server.

Creating and using a key database or key ring

- The LDAP client and server use the System SSL functions provided in z/OS to set up SSL/TLS protected communications. The System SSL capability requires a key database or key ring to be set up before SSL/TLS protected communications can begin.
- The key database is a password protected file stored in the hierarchical file system (HFS). This file is created and managed using a utility program provided with System SSL called **gskkyman**. Directions for using the **gskkyman** utility can be found in z/OS: *System Secure Sockets Layer Programming*. The key database file that is created must be accessible by the LDAP server.
- The key ring is maintained by RACF. This object is created and managed using the RACF Digital Certificate command, **RACDCERT**. Directions for using the **RACDCERT** command can be found in *z/OS:*Security Server RACF Command Language Reference.
- The user ID under which the LDAP server runs must be authorized by RACF to use RACF key rings. To authorize the LDAP server, you can use the RACF commands in the following example:

RDEFINE FACILITY IRR.DIGTCERT.LIST UACC(NONE)
RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(NONE)
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(LDAPSRV) ACCESS(CONTROL)
PERMIT IRR.DIGTCERT.LIST CLASS(FACILITY) ID(LDAPSRV) ACCESS(CONTROL)

- Remember to refresh RACF after doing the authorizations.
- | SETROPTS RACLIST(FACILITY) REFRESH
- For testing purposes, the LDAP server can use a self-signed certificate. In this case, the certificate of the LDAP server must also be stored in the key database, or key ring, of the LDAP client in order for SSL/TLS protected LDAP communications to work between the client and server.
- The certificate that the LDAP server is going to use should be stored as the default certificate in a key database or key ring, or the certificate label of the certificate is configured in the LDAP server using the **sslCertificate** option.

Obtaining a certificate

The LDAP server or client can obtain a certificate by contacting a certificate authority (CA) and requesting a certificate. Utilities to formulate a certificate request are provided by System SSL, **gskkyman**, and RACF, **RACDCERT**. This certificate request is usually passed to the CA by means of an electronic mail message or by an HTML form which is filled out using a web browser. Once the CA verifies the information for the LDAP client or server, a certificate is returned to the requester, usually by an electronic mail message. The contents of the mail message are used to define the certificate in the key database or key ring.

Setting up the security options for the LDAP server

The following options for SSL/TLS can be set in the **slapd.conf** file. They are described in detail in "Configuration file options" on page 53.

- listen
- sslAuth
- sslCertificate
- sslCipherSpecs
- sslKeyRingFile
- sslKeyRingFilePW
- sslKeyRingPWStashFile

Notes:

1. The **replKeyRingFile** and **replKeyRingPW** options are no longer necessary or recognized by the LDAP server. These options should be removed from the configuration file.

	information, see the listen option on page 59.
	 LDAP can be configured for SSL/TLS in two ways: For secure only communication, specify one or more listen options for secure communications in the following format: <pre> ldaps://[IP_address hostname] [:portNumber]</pre> For bimodal (non-secure/secure) communication, specify one or more listen options for non-secure communications in the following format: <pre> ldap://[IP_address hostname] [:portNumber]</pre> For more information on the listen option, refer to page 59.
	sslKeyRingFile specifies the name of the key database or the key ring used by the LDAP server. This key database or key ring is also used for SSL/TLS protected replication. Because the replicating server may be acting as both a replica server and an LDAP server, the replica server's certificate (or CA's certificate) must be contained in the replicating server's key database file or key ring.
	A key database requires a password. The password may be specified on the sslKeyRingFilePW option or the name of a password stash file may be specified on the sslKeyRingPWStashFile option in the configuration file. Use of a stash file provides a method of specifying a password in a form that can not be easily read by a human. The gskkyman utility provides a function to create the key database password stash file.
	When a RACF key ring is used instead of a key database, the sslKeyRingFilePW and sslKeyRingPWStashFile should not be specified in the configuration file.
	The LDAP server is configured to provide server and, optionally, client authentication. The sslAuth option controls this setting.
	With server authentication, the LDAP server must have a digital certificate (based on the X.509 standard). This digital certificate is used to authenticate the LDAP server to the LDAP client application. The LDAP server supplies the client with the LDAP server's X.509 certificate during the initial SSL handshake. If the client validates the server's certificate, then a secure, encrypted communication channel is established between the LDAP server and the LDAP client application.
	In addition, if the LDAP server is configured to use server and client authentication, and the client sends a digital certificate on the initial SSL handshake, it must be validated by the LDAP server before the secure encrypted communication channel is established between them.
	Client authentication by the LDAP server facilitates the use of certificate bind (SASL mechanism of EXTERNAL) by an LDAP client. The bind identity becomes the distinguished name in the client digital certificate.
	Note: If the LDAP server is configured for both server and client authentication, but a client does not send a digital certificate, then the server will act as if configured for server authentication only. This provides backward compatibility of the LDAP server.
	The sslCertificate option indicates the label of the server certificate that is to be used. If the default certificate has not been set in the key database or key ring, or if a certificate other than the default certificate is desired then this option is needed.
	The sslCipherSpecs option specifies the cipher specifications that will be accepted from clients. If this option is not specified then all cipher specifications supported by the LDAP server will be used. Depending upon the level of System SSL support, the list of acceptable cipher specifications may be lowered because certain specifications may not be supported by System SSL for that level of the product.

Setting up an LDAP client

As with the LDAP server, the LDAP client that wishes to use SSL/TLS protected communication needs access to a key database or key ring. If the LDAP server you are going to contact is using a self-signed certificate (as is done frequently while testing SSL/TLS protected communications between an LDAP client and server), then the self-signed certificate of the LDAP server must be stored into the LDAP clients key database or key ring.

If the LDAP server you are going to contact is using a certificate which is signed by a certificate authority (CA), you must ensure that the certificate for the CA is contained in the key database. Use whatever means is provided by the CA for obtaining the CA certificate. The certificate should be obtainable in a format that is acceptable to the gskkyman utility or RACDCERT command.

If the LDAP server is configured for server and client authentication and the client wants client authentication to occur, then the LDAP client must obtain its own certificate from a CA and store it in the clients own key ring or key database and mark it as the default.

Once the key database file or key ring is created and contains the proper certificates, then the LDAP client is ready to perform SSL/TLS protected communications with an LDAP server. The LDAP Operation Utilities (for example, Idapsearch) can be used to communicate securely with the LDAP server using a secure only connection. The utilities are explained in z/OS: Security Server LDAP Client Programming.

Using LDAP client APIs to access LDAP using SSL/TLS

The Idap ssl client init and Idap ssl init APIs can be used to start a secure only connection to an LDAP server. A description of these APIs can be found in z/OS: Security Server LDAP Client Programming.

Support of certificate bind

The SASL bind mechanism of EXTERNAL is supported by the LDAP server. This means that the authentication on the bind is performed using the data obtained during the SSL/TLS client authentication that was performed on the SSL/TLS handshake with the client.

To use SASL External bind, the following steps must occur:

- . The LDAP server must be configured and started with sslAuth set to serverClientAuth so that the server can authenticate the client.
- · The client connects to the LDAP server and performs the SSL/TLS handshake. The handshake sends the client certificate to the LDAP server.
- The client performs a SASL bind with the mechanism of EXTERNAL.

At this point, the LDAP server will consider the bind DN of the client for authorization purposes to be the client's DN as transmitted in the client's certificate on the handshake.

Configuring for user password encryption

The LDAP server allows prevention of unauthorized access to user passwords in the TDBM backend. The userPassword attribute values can be encoded when stored in the directory, which prevents clear text passwords from being accessed by any users, including the system administrators. In the current implementation, only the userPassword attribute values are encrypted. Use of the terms "user password" and "password" refer to the userPassword attribute. Use of the term "user entry" refers to an entry in TDBM that contains a userPassword attribute.

Note: The z/OS LDAP server does not allow userPassword attributes in distinguished names.

The administrator may configure the server to encode userPassword attribute values in either a one-way hash format or a two-way, symmetric, encryption format.

After the server is configured and started, any new passwords for new user entries, or modified passwords for existing user entries are encoded before they are stored in the TDBM backend. The encoded passwords are tagged with the encoding algorithm name so that passwords encoded in different formats can coexist in the directory. When the encoding configuration is changed, existing encoded passwords remain unchanged and continue to be usable.

When Idif2tdbm is used to load a TDBM backend, all clear text user passwords in new entries are encrypted by the method specified in the configuration file.

If there are encrypted userPassword values in the LDAP database, the unload utility tdbm2ldif unloads the TDBM backend to LDIF format with the password in the binary format of:

```
userPassword:: base64encoded and tag encryptedvalue
```

This format can be loaded by the load utility, for example at a replica server, and preserves the encrypted passwords.

Note about password encryption and replication: Some important considerations are described in "Password encryption and replication" on page 269.

The -t option on tdbm2ldif unloads the user passwords in an encrypted format that might be more appropriate for loading by non-z/OS LDAP servers. The LDIF format unloaded by the -t option can be called encryption "tag visible" format and looks like:

```
userPassword: \{tag\}base64encoded and encryptedvalue
```

where tag is none, crypt, MD5, SHA, or DES:keylabel.

In this format the tag is visible, and only the password itself is encrypted and base64 encoded.

Notes:

- 1. The tag is enclosed in a left brace and a right brace. One colon is used between the userPassword keyword and the value, as opposed to two colons in the standard LDIF format of userPassword unloaded by tdbm2ldif. The format produced by tdbm2ldif -t cannot be read by the z/OS ldif2tdbm utility programs. It is intended for other LDAP platforms or tools that may be able to interpret this LDIF format with the encryption tag visible.
- 2. The values returned by the crypt() algorithm are not portable to other X/Open-conformant systems. This means that user password values encoded by the crypt() algorithm and unloaded as tagged output using tdbm2ldif -t are not portable when loaded by another platform's load utility.

If the TDBM backend is already loaded and the LDAP server is running, the db2pwden utility is provided to encrypt all clear text userPassword attribute values in the method configured on the LDAP server.

The db2pwden utility is similar to the LDAP operation utilities, such as Idapsearch, in that it acts like a client to the LDAP server and has similar command line options. See "db2pwden utility" on page 148 for information about the db2pwden utility and see z/OS: Security Server LDAP Client Programming for more information about LDAP operation utilities, such as Idapsearch. The db2pwden utility must be run by the LDAP server administrator using the **adminDN** and password configured on the server.

Be aware that once a password is encrypted in a one-way hash, its clear text value can no longer be retrieved or displayed.

One-way hash formats:

crypt

MD5

SHA

A crypt, MD5, or SHA hashed password can be used for password matching on an LDAP simple bind, but it cannot be decrypted. During simple bind, the bind password is hashed and compared with the stored userPassword attribute values for matching verification.

MD5 and SHA hashing require the z/OS Open Cryptographic Services Facility (OCSF) be installed and configured, and the necessary security authorizations to be set up for the LDAP server user ID in RACF. See the configuring information in z/OS: Open Cryptographic Services Facility Application Programming for instructions on how to do this.

For applications which require retrieval of clear passwords, such as middle-tier authentication agents, the directory administrator must configure the LDAP server to perform either a two-way encoding or no encryption of user passwords. Access to password data stored in the directory can be protected by the access control mechanism of the directory.

Two-way encryption format:

DES

The DES algorithm is provided to allow values of the userPassword attribute to be encoded in the TDBM backend and retrieved as part of an entry in the original clear format. Some applications such as middle-tier authentication servers require passwords to be retrieved in clear text, however, corporate security policies might prohibit storing clear passwords in a secondary permanent storage. This option satisfies both requirements.

A DES encrypted password can be used for password matching on a simple bind and can be decrypted to be returned as clear text on a search request when the client is authorized to do so. During simple bind, the bind password is encrypted and compared with the stored version for matching verification. During a search, if the client is authorized through directory access controls to see the userPassword attribute value, then it is decrypted and returned as clear text.

DES encryption requires OCSF be installed and configured, and the necessary security authorizations to be set up for the LDAP server user ID in RACF. See the configuring information in z/OS: Open Cryptographic Services Facility Application Programming for instructions on how to do this.

The DES algorithm also requires a key label and the single-length data key it refers to for password matching and decryption to take place. When a userPassword is stored in TDBM, the key label is stored along with the tag and encrypted user password. The Integrated Cryptographic Service Facility (ICSF), Key Generator Utility Program (KGUP) and Cryptographic Key Data Set (CKDS) are used to generate and store the key label and RACF is used to limit access to this DES encryption key to only the LDAP server. See the information on managing cryptographic keys and using the Key Generator Utility Program in z/OS: ICSF Administrator's Guide for information on generating, storing into CKDS, and authorizing a Data-Encrypting Key. Remember to refresh CKDS and RACF after entering and authorizing a key.

The DES key label must correspond to the same DES keys across a sysplex and be accessible to all LDAP servers that are using the same TDBM backend. It is recommended that you use one DES key label. If multiple DES key labels are used by different servers in the sysplex, for example, then all the servers in the sysplex need to have access to all the keys.

A simple bind will succeed if the password provided in the bind request matches with any of the multiple values of the userPassword attribute. Note that depending on when userPassword values are stored in the directory, different attribute values can be encoded using different encoding methods.

Note: The crypt() algorithm, implemented across many platforms, accepts only the first eight characters of a password. As a result, any password supplied on an Idap simple bind operation or Idap compare operation that matches the first eight characters of a userPassword attribute value encrypted with the crypt() algorithm in the directory will match.

Chapter 8. Customizing the LDAP server configuration

This chapter contains information on how to set up the **slapd.conf** configuration file and how to configure the LDAP server to run with the options you choose. The **slapd.conf** file is also used by the LDAP utilities.

- · "Creating the slapd.conf file"
- "Configuration file options" on page 53
- · "Configuration considerations" on page 74
- "Determining operational mode" on page 75
- "Establishing the administrator DN and password" on page 80
- "Example configuration scenarios" on page 82

Creating the slapd.conf file

This section discusses what is necessary for creating the **slapd.conf** configuration file. Specifically, this section:

- · Describes where the slapd.conf file is located
- · Shows the configuration file format
- · Provides a checklist for the configuration file options
- · Lists all of the configuration file options
- · Describes how to establish the administrator DN and password
- Discusses encryption using the **pwEncryption** option

Appendix A, "LDAP server configuration file (slapd.conf)" on page 371 shows the configuration file that is provided for you.

Locating slapd.conf

All LDAP server runtime configuration is accomplished through the configuration file **slapd.conf**, installed in the **/usr/lpp/ldap/etc** directory. If this is your first time installing the LDAP server, create a new copy of **slapd.conf** with:

cp /usr/lpp/ldap/etc/slapd.conf /etc/ldap/slapd.conf

and edit /etc/ldap/slapd.conf.

An alternate configuration file can be specified through a command-line option to the LDAP server and other LDAP programs.

The initial configuration contains default versions of some configuration settings. It does not contain a database suffix.

Configuration file format

The **slapd.conf** file consists of the following sections:

Global section

Contains configuration options that apply to the LDAP server as a whole (including all backends).

SDBM backend-specific section

Contains configuration options that apply to the SDBM backend.

TDBM backend-specific section

Contains configuration options that apply to the TDBM backend. It is possible to have one or more of these sections depending on how many TDBM backends your installation uses.

EXOP backend-specific section

Contains only the **database** statement necessary for the EXOP backend.

Noted below are some rules for setting up **slapd.conf**:

- The configuration file always contains a global section followed by one or more database backend sections that contain information specific to a backend instance.
- The configuration file or files must be in code page IBM-1047.
- For single-valued options that appear more than once, the last appearance in the slapd.conf file is used.
- Blank lines and comment lines beginning with the pound sign character (#) are ignored.
- If a line begins with one or more blank spaces, it is considered a continuation of the previous line.
- · If an argument contains one or more blank spaces, the argument should be enclosed in double quotation marks (for example, "argument one"). If an argument contains a double quotation mark or a backslash character (\), the double quotation mark or backslash character should be preceded by a backslash character (\).
- A pound sign (#) cannot be used at the end of a configuration line to denote commentary.

Figure 6 shows the general format of slapd.conf.

```
# Global options - these options apply to every database
<global configuration options>
# SDBM database definition and configuration options
database sdbm <GLDBSDBM>
<configuration options specific to SDBM backend>
# TDBM database definition and configuration options
database tdbm <GLDBTDBM>
<configuration options specific to TDBM backend>
# EXOP database definition and configuration options
database exop <GLDXPDIR>
```

Figure 6. General format of slapd.conf

Configuration file checklist

The following table is provided to assist you in determining which configuration file options you will need to use in your slapd.conf file. Depending on the section in the configuration file (Global, SDBM, TDBM, or EXOP), certain topics (SSL, schema, replication, and so on) have options that are required or optional.

Table 9. Configuration file options checklist

Section/topic	Check	Options	
Global		adminDN is required	
		adminPW, altServer, commThreads, digestRealm, idleConnectionTimeout, include, listen, logfile, maxConnections, pcThreads, referral, sendV3stringsoverV2as, serverEtherAddr, sizeLimit, timeLimit, and validateincomingV2strings are optional.	
SSL/TLS		sslAuth, sslCertificate, sslKeyRingFilePW, sslKeyRingPWStashFile, and sslCipherSpecs are optional	
		sslKeyRingFile is required if a listen option is initialized for secure socket communications or a listen option is initialized for non-secure socket communications that is intended to support switching to secure socket communications once the connection is established.	

Table 9. Configuration file options checklist (continued)

Section/topic	Check	Options	
Sysplex		sysplexGroupName and sysplexServerName are optional	
Kerberos		supportKrb5 and serverKrbPrinc are required	
		krbKeytab and krbLDAPAdmin are optional	
SDBM backend		database and suffix are required	
		sizeLimit and timeLimit are optional	
Kerberos		krbldentityMap is optional	
TDBM backend		database, databasename, dbuserid, servername, and suff are required	
		attrOverflowSize, dsnaoini, extendedGroupSearching, readonly, sizeLimit, and timeLimit are optional	
Password encryption		pwEncryption is optional	
Replication		masterServer, masterServerDN, and masterServerPW are optional	
Multi-server		multiserver is optional	
Kerberos		krbldentityMap is optional	
Native authentication		useNativeAuth and nativeAuthSubtree are required	
		nativeUpdateAllowed is optional	
EXOP backend		database is required	

Note: Be sure to specify **adminDN**. You can specify the **adminPW** here or in a database entry. See "Establishing the administrator DN and password" on page 80 for more information. Note that the use of the **adminPW** option is strongly discouraged. Instead, an existing entry in the directory should be designated as the **adminDN**.

Configuration file options

This section contains an alphabetical listing of the configuration file options. For each option, a table shows an **X** in the areas (Global, TDBM, SDBM, and EXOP) of the configuration file where the option can be used.

Specifying a value for filename

In the configuration file options, the value for filename can be specified in one of the following ways:

/pathname/filename

Specifies the full path name of a file in the USS Hierarchical File System (HFS).

filename

Specifies a path name that is relative to the current working directory of the LDAP server. Note that when running from a started task or batch, there is no current working directory defined. This format is not recommended.

//'dataset.name'

Specifies the fully-qualified name of a configuration file stored in a sequential dataset.

//'dataset.name(member)'

Specifies the fully-qualified name of a configuration file stored in a partitioned dataset.

//DD:DDNAME

Specifies the DDNAME of a configuration that has been specified as a DD card in the JCL for the batch job or started task.

adminDN dn

Global	TDBM	SDBM	EXOP
X			

The distinguished name (DN) of the administrator for this LDAP server. This DN will have unrestricted access to all entries in the local directory. The name that is chosen should be descriptive of the person that will have knowledge of and administer the LDAP server. The format of the name must be in DN format which is described in Chapter 15, "Data model" on page 155. It is recommended, though not necessary, that the DN have the same suffix as one of the suffix option values in the configuration file.

"Establishing the administrator DN and password" on page 80 describes how to set up your administrator DN.

With LDAP V3 support, UTF-8 characters can be used for textual attributes stored in the directory. It is also desirable to allow any UTF-8 character to appear in distinguished names, and in particular, the adminDN distinguished name. Because the LDAP configuration files are defined to hold information only in the IBM-1047 character set, a solution is required for the configuration files that allows you to use distinguished names containing UTF-8 characters but using only the IBM-1047 character set. To solve this problem, an escape mechanism has been introduced for purposes of entering either the adminDN or the masterServerDN. This escape mechanism allows the entry of UTF-8 characters while keeping the input string value to within the IBM-1047 character set. The escape mechanism employed requires that you express UTF-8 characters which are not within the X'00' - X'7F' range (7-bit ASCII which is the single-byte form of UTF-8 characters) in the form of a set of four character representations. This representation has the form "&nmm" where 0<n<3 and 0<m<7. You might recognize nmm as being an octal value for a byte of information. Thus, if you want to create an **adminDN** which was the following distinguished name:

cn=Peter <U umlaut>nger, o=Widgets, c=DE

enter the adminDN into this option value as:

cn=Peter &nmm&nmmnger, o=Widgets, c=DE

Because the <U umlaut> is not within the 7-bit ASCII range, the value must be escaped to the octal representation of the UTF-8 multi-byte character. In the case of <U umlaut>, the Unicode code point is X'00DC'. Converted to UTF-8, this character is a multi-byte sequence: X'C3BC'. (Refer to "UTF-8 support" on page 151 for conversion information.) Converted to the escaped form for input into the adminDN field, this character is represented as "&303&234" since X'C3' is octal 303 and X'BC' is octal 234. Thus, the adminDN above would be entered as:

cn=Peter &303&234nger, o=Widgets, c=DE

If there is a case where you need to enter an adminDN string which contains the string "&nmm" where 0 < n < 3 and 0 < m < 7, then you must escape the ampersand by using its octal representation which is "&046".

adminPW string

Global	TDBM	SDBM	EXOP
X			

The password of the administrator (adminDN) for this server.

"Establishing the administrator DN and password" on page 80 describes how to set up your administrator password.

Note: Use of the **adminPW** configuration option is strongly discouraged in production environments. Instead, specify your **adminDN** as the distinguished name of an existing entry in the directory information tree. This will eliminate passwords from the configuration file.

altServer Idap URL

Global	TDBM	SDBM	EXOP
X			

Specifies an equivalent server to this LDAP server. It may or may not be a replica, but should contain the same naming contexts. The alternate server is specified as an LDAP URL.

In the following example, myldap.server.com is the host name and 3389 is the port number of the LDAP URL:

altServer ldap://myldap.server.com:3389

Use the **altServer** configuration option to list alternate servers in the **altServer** root DSE attribute. The **altServer** root DSE attribute is used to list other servers that may be contacted in case this server is not available at some later time.

attribute name [name2 ...namen] {bin | ces | cis | tel | dn} colname maxlen {normal | sensitive | critical}

Global	TDBM	SDBM	EXOP
X			

Note: The **attribute** option is now ignored by the LDAP server.

attrOverflowSize num-of-bytes

Global	TDBM	SDBM	EXOP
	X		

Specifies, in bytes, the minimum size of an attribute value required to store the value in a long attribute value table. The choice of this value allows large attribute values (such as **JPEG** and **GIF** files) to be stored in a separate DB2 table in a separate DB2 table space. The maximum size of this value is 2147483647.

Default = 255

commThreads int

Global	TDBM	SDBM	EXOP
Х			

Specifies the number of threads to be initialized for the communication thread pool. This thread pool handles the connections between the LDAP server and its clients.

Default = 10

It is recommended that **commThreads** be set to approximately two times the number of CPUs that are running in your LPAR. However, this is a general rule depending upon the activity that your LDAP server experiences.

The commThreads option deprecates the maxThreads and waitingThreads options, which are no longer evaluated by the LDAP server. (See "Migrating from maxThreads and waitingThreads to commThreads option" on page 107 for migration information.)

database dbtype dblibpath [name]

Global	TDBM	SDBM	EXOP
	X	X	X

Marks the beginning of a new database section.

- For dbtvpe:
 - IBM supports tdbm (DB2), sdbm (RACF), and exop (extended operations). The type config is reserved by the LDAP server and should not appear as **dbtype** in your configuration files.
- - This is the file name of the shared library (DLL) containing the backend database code. Unless you have changed the names of the LDAP DLLs, specify GLDBTDBM when dbtype is tdbm, GLDBSDBM when dbtype is sdbm, and GLDXPDIR when dbtype is exop.
- For name:
 - This optional value is a name that is used to identify this backend. If specified, the *name* must be different (ignoring case) than the name specified on any other database record in this configuration file. You cannot specify **cdbm1** as the *name*.

databasename dbname

	Global	TDBM	SDBM	EXOP
1		X		

Specifies the name of the DB2 database this backend uses to store directory data.

dbuserid userid

Global	TDBM	SDBM	EXOP
	X		

Specifies a z/OS user ID that will be the owner of the DB2 tables.

digestRealm hostname

Global	TDBM	SDBM	EXOP
X			

Specifies a realm name to be used when doing DIGEST-MD5 or CRAM-MD5 SASL authentication binds to the LDAP server. The digestRealm is used to help calculate a hash for DIGEST-MD5 and CRAM-MD5 authentication binds. It is suggested that the hostname be a DNS-host name and not an IP address.

Default = primary host name of the server obtained from DNS.

dsnaoini filename

Global	TDBM	SDBM	EXOP
	X		

Specifies the name of the CLI Initialization file or sequential data set (or PDS member) you created in step 4 on page 14. If the **dsnaoini** option is set in the configuration file, the LDAP server will export the **DSNAOINI** environment variable to the value specified for the configuration option.

In addition to using the **dsnaoini** configuration option or the **DSNAOINI** environment variable, a DSNAOINI DD card can be used to specify the CLI Initialization file. If the DSNAOINI DD card is specified in the JCL for the job/started task for the LDAP server, then neither the **dsnaoini** configuration option value nor the **DSNAOINI** environment variable need to be specified. "Running the LDAP server using data sets" on page 96 gives more information on this process. See *DB2 for OS/390 and z/OS: ODBC Guide and Reference* for details on ways to specify the CLI initialization file.

See *DB2* for *OS/390* and *z/OS*: *ODBC* Guide and Reference for other ways in which the DSNAOINI ODBC initialization file can be specified. In order for the TDBM backend to run, the initialization file must be specified in one of the ways indicated.

Refer to page 53 for information on specifying the *filename*.

extendedGroupSearching {on | off}

Global	TDBM	SDBM	EXOP
	Х		

Specifies whether a backend participates in extended group membership searching on a client bind request. If this option is **on**, group memberships are gathered from this backend during LDAP bind processing in addition to the backend in which the bind DN resides. If this option is **off**, group memberships are not gathered from this backend unless the bind DN resides in this backend. The default is **off**.

The group memberships gathered on a client's LDAP bind request are used for authorization checking of the client's request. The administrator should know, in general, which backends may contain group information so they can be marked for **extendedGroupSearching**. Group memberships are necessary for complete authorization checking of a client request.

The server control **authenticateOnly** is supported by the LDAP server so that a client can override both **extendedGroupSearching** and group membership gathering from the backend where the DN resides. See Appendix G, "Supported server controls" on page 431 for more information.

This option applies only to the backend in which it is defined.

idleConnectionTimeout int

Global	TDBM	SDBM	EXOP
Х			

Specifies the amount of time in seconds that the LDAP server will wait on an idle connection before timing out on a particular client connection.

Default = 0 (indefinitely)

Recommended value = 1800 (30 minutes) Minimum value = 30 seconds

include filename

0	Blobal	TDBM	SDBM	EXOP
	Х	X	X	X

Specifies the path and file name of a file to be included as a part of the LDAP server configuration.

Refer to page 53 for information on specifying *filename*.

Note that the LDAP server will not detect loop conditions in a set of included files. Configuration may encounter errors or fail if the same file is processed more than once. While nested include files are supported, including the same file in such a way as to form a loop condition is not supported.

index attrlist [eq | none]

Global	TDBM	SDBM	EXOP

Note: The **index** option is now ignored by the LDAP server.

krbldentityMap [on | off]

Global	TDBM	SDBM	EXOP
	X	X	

Specifies if this backend will participate in Kerberos identity mapping. If it will participate, then the server will attempt to map the Kerberos identity that performed the bind to DNs that exist in this backend. The mapped DNs will then be used for access control.

Default = off

krbKeytab {serverKeytab | none}

Global	TDBM	SDBM	EXOP
X			

Specifies the keytab that will be used for the LDAP server. Following is an example: krbKeytab /home/users/u1/keytab

The value for this option must be specified if the KDC does not exist on the same machine as the LDAP server. However, if the KDC resides on the same machine as the LDAP server and the user ID the LDAP server runs under has permission to read from the IRR.RUSERMAP facility class in RACF, then the krbKeytab value should be set to none.

krbLDAPAdmin kerberosldentityDN

Global	TDBM	SDBM	EXOP
Х			

Specifies the Kerberos identity that represents the LDAP administrator. This option allows the administrator to bind through Kerberos and still maintain administrative authority. The value for this option must be specified as a DN with the attribute type of **ibm-kn**. Following is an example:

krbLDAPAdmin ibm-kn=LDAPAdmin@myrealm.com

listen Idap_URL

Global	TDBM	SDBM	EXOP
X			

Specifies, in LDAP URL format, the IP address (or host name) and the port number where the LDAP server will listen to incoming client requests. This parameter may be specified more than once in the configuration file.

Note that the **listen** value may be established in the configuration file, or it may be established using the optional start-up parameter for **listen** (see "Setting up and running the LDAP server in the z/OS shell" on page 97).

Default = INADDR_ANY (that is, Idap://:389) on the nonsecure default port 389.

The format of *Idap_URL* for the **Iisten** option to listen on a TCP/IP socket interface is the following:

```
{ldap:// | ldaps://}[IP_address | hostname][:portNumber]
```

The format of *ldap_URL* for the **listen** option to listen on the z/OS SAF interface is the following: {1dap:// | 1daps://}:pc

where:

Idap:// Specifies that the server listen on nonsecure addresses or ports. Note that if SSL/TLS is configured for the server, then once a connection is established, the client may switch to secure communication using the Start TLS Extended Operation.

Idaps://

Used to have the server listen on secure addresses or ports. Once a connection is established to the server, the client must begin the SSL/TLS handshake protocol.

IP address

Specifies the IP address.

hostname

Specifies the host name. If the host name is used for the **listen** option, all of the IP addresses are obtained and the LDAP server listens on each of these IP addresses.

portNumber

Specifies the port number. The *portNumber* is optional. If the port number is not specified for an **Idap://**, then the default of 389 is used for nonsecure connections. If the port number is not specified for an **Idaps://**, then the default of 636 is used for secure connections.

Range = 1-65536

If the **sysplexGroupName** and **sysplexServerName** options are present in the configuration file, the port number specified for this server instance must be the same as the port number specified for all other members of the same *group_name* in the sysplex for dynamic workload balancing to function properly.

It is advisable to reserve the port number or numbers chosen here in your TCP/IP profile data set. Also, be aware that port numbers below 1024 may require additional specifications. Consult z/OS: Communications Server: IP Configuration Reference for more information.

рс Specifies that the LDAP server should listen for program call (PC) calls from Policy Director using the z/OS Security Authorization Facility (SAF) interface.

Note that when the listen option is initialized to listen for PC calls on the LDAP server, the listen parameter must not include an IP address or a host name.

Also, there is no difference if you specify **Idap** or **Idaps** as part of *Idap_URL*. Both produce the same result.

Following are some examples of how you can specify *Idap_URL*.

If you specify:

1dap://

the LDAP server binds and listens on all available IP addresses (INADDR ANY) on the system on the nonsecure default port of 389.

If you specify:

ldap://us.endicott.ibm.com:489

the LDAP server binds and listens on all of the IP addresses associated with host name us.endicott.ibm.com on the nonsecure port of 489 for incoming client requests.

· If you specify:

ldap://9.130.77.27

the LDAP server binds and listens on IP address 9.130.77.27 on the default nonsecure port of 389 for incoming client requests.

If you specify:

ldaps://us.endicott.ibm.com

the LDAP server binds and listens for incoming client requests on the IP address or addresses associated with host name us.endicott.ibm.com on the default secure port of 636.

· If you specify:

ldaps://9.130.77.27:736

the LDAP server binds and listens on IP address 9.130.77.27 on the secure port of 736.

If you specify:

1dap://:489

the LDAP server binds and listens on all available IP addresses (INADDR_ANY) on the system on the nonsecure port of 489 for incoming client requests

· If you specify:

1daps://:777

the LDAP server binds and listens on all available IP addresses (INADDR ANY) on the system on the secure port of 777 for incoming client requests.

If you specify:

ldap://:pc

the LDAP server binds and listens for PC calls from Policy Director using the SAF interface into the server.

Note: The listen parameter deprecates the security, port, and securePort options in the configuration file. If there is a listen option specified in the configuration file along with either security, port, or securePort, the listen option takes precedence over what has been specified for security, port, or securePort. If using an earlier version of the configuration file with security, port, or securePort, the LDAP server will be configured to listen on the port numbers specified for securePort, port, or both depending upon the security setting. However, it is highly recommended that the LDAP server be configured using the listen option.

To migrate from one or more of the port, securePort, or security options to the listen configuration option, refer to "Migrating from port, securePort, and security to listen option" on page 106.

logfile filename

Global	TDBM	SDBM	EXOP
X			

Specifies where to place the activity log records when activity logging is enabled.

Refer to page 53 for information on specifying the *filename*.

masterServer Idap URL

Global	TDBM	SDBM	EXOP
	X		

Specifies the location of this replica's master server in LDAP URL format.

In the following example, myldap.server.com is the host name and 3389 is the port number of the LDAP URL:

masterServer ldap://myldap.server.com:3389

masterServerDN dn

Global	TDBM	SDBM	EXOP
	Х		

Specifies the DN allowed to make changes to the replica. The format of the name must be in DN format which is described in Chapter 15, "Data model" on page 155. The presence of this option indicates that the server instance using this configuration file is a slave.

In order to enter characters in this distinguished name which are outside of the 7-bit ASCII range when expressed in Unicode (or UTF-8), then you must escape these characters. See the description under the adminDN configuration option (page 54) for details on how to do this.

masterServerPW string

Global	TDBM	SDBM	EXOP
	Х		

Specifies the password for the masterServerDN that will be allowed to make updates. This option is only applicable for a slave LDAP server. See "Establishing the administrator DN and password" on page 80 for additional information about the master server password.

Note: Use of the masterServerPW configuration option is strongly discouraged in production environments. Instead, specify your masterServerDN as the distinguished name of an existing entry in the directory information tree. This will eliminate passwords from the configuration file.

maxConnections int

Global	TDBM	SDBM	EXOP
X			

Specifies the maximum number of concurrently connected clients that the LDAP server allows. The number is specified as a positive integer.

Default = operating system maximum

Range = 30 to 65535

The LDAP server limits the number of client connections by restricting the number of file and socket descriptors used by the LDAP server. Some of the descriptors are used by the LDAP server for its own file descriptors and passive socket descriptors. The value specified for this option should take into account that the server uses approximately 10 descriptors for internal functions and will use more depending upon the number of additional sockets used as passive sockets for connection attempts by clients.

The maximum number of client connections is further restricted by:

- · The maximum number of files a single process can have concurrently active. The MAXFILEPROC statement for BPXPRMxx and the maxfileproc option on the RACF altuser command are used to set the limit. Only processes with superuser authority can adjust the limit beyond the limit specified by MAXFILEPROC. Attempts to exceed this limit by non-superuser processes may be audited by the security manager.
- The maximum number of sockets allowed by the TCP/IP socket file system. The MAXSOCKETS option on the NETWORK statement for BPXPRMxx sets this limit.

Setting these limits too high can affect system performance by using too many resources and deprive other functions of their share of the same resources.

Note: If using a configuration file from z/OS LDAP Release 1 or earlier, it may be necessary to update the maxConnections value. Refer to "Migrating to the updated maxConnections option" on page 108.

maxThreads int

Global	TDBM	SDBM	EXOP
Х			

Note: The maxThreads option is now ignored by the LDAP server. Use the commThreads option to specify the number of threads that will be created to handle incoming work from the clients. Refer to page 55 for a description of the **commThreads** option.)

multiserver {Y | y | N | n}

Global	TDBM	SDBM	EXOP
	Х		

Indicates the operating mode in which this server will run. Specifying either **y** or **Y** indicates the server runs in multi-server mode with or without dynamic workload management enabled (see page 75 for a description of multi-server operating modes). Specifying either **n** or **N** indicates the server runs in single-server mode.

If **n** or **N** is specified, and both **sysplexGroupName** and **sysplexServerName** options are present in the configuration file, the **multiserver** option value is overridden and the server operates in multi-server mode.

The multiserver keyword may be present without the sysplexGroupName and sysplexServerName keywords, in which case replication is disabled in the server and sysplex Workload Management features are also disabled.

If **sysplexGroupName**, **sysplexServerName**, and **multiserver** keywords are all omitted from the server configuration file, the server will operate in single-server mode and replication will be enabled.

nativeAuthSubtree {all | dn}

Global	TDBM	SDBM	EXOP
	X		

Specifies the distinguished name of a subtree where all of its entries are eligible to participate in native authentication. This parameter can appear zero or more times to specify all subtrees that use native authentication. If this parameter is omitted, contains no value, or is set to **all**, then the entire directory is subject to native authentication. This option is only valid if the **useNativeAuth** option is also specified.

nativeUpdateAllowed {on | yes | off | no}

Global	TDBM	SDBM	EXOP
	X		

Enables native password changes in the Security Server to occur through the TDBM backend if the **useNativeAuth** option is specified.

Default = off

objectclass name [requires attrs] [allow attrs]

Global	TDBM	SDBM	EXOP
Х			

Note: The objectclass option is now ignored by the LDAP server.

pcThreads int

Global	TDBM	SDBM	EXOP
X			

Specifies the number of threads to be initialized to handle incoming program call (PC) calls using the z/OS SAF interface into the LDAP server.

Default = 10

port int

Global	TDBM	SDBM	EXOP
Х			

Note: The port option is deprecated when the listen option is specified. See page 59 for information on the listen option.

TCP/IP port used for non-SSL communications.

Default = 389Range = 1 - 65535

If the sysplexGroupName and sysplexServerName options are present in the configuration file, the port number specified for this server instance must be the same as the port number specified for all other members of the same group_name in the sysplex for dynamic workload balancing to function properly. Note that the port number may be established in the configuration file, or it may be established using the optional startup parameter for port (see "Setting up and running the LDAP server in the z/OS shell" on page 97).

It is advisable to reserve the port number chosen here in your TCP/IP profile data set. Consult z/OS: Communications Server: IP Configuration Guide for further information.

pwEncryption {none | crypt | MD5 | SHA | DES: keylabel}

Global	TDBM	SDBM	EXOP
	Х		

Specifies what encryption method to use when storing the userPassword attribute values in the backend of the directory.

- Specifies no encryption. Specifies the **userPassword** attribute values are stored in clear text format. Note that these clear text passwords are stored in the directory prefixed with the tag {none}.
- Specifies that userPassword attribute values are encoded by the crypt() algorithm before they are stored in the directory. These passwords are stored in the directory prefixed with the tag {crypt}.
- MD5 Specifies that userPassword attribute values are encoded by the MD5 hash algorithm using OCSF before they are stored in the directory. These passwords are stored in the directory prefixed with the tag {MD5}.
- SHA Specifies that userPassword attribute values are encoded by the SHA hash algorithm using OCSF before they are stored in the directory. These passwords are stored in the directory prefixed with the tag {SHA}.

DES: keylabel

Specifies that userPassword attribute values are encrypted by the DES algorithm using the specified key label using OCSF and ICSF before they are stored in the directory, and can be retrieved as part of an entry in the original clear text format. These passwords stored in the directory are prefixed with the tag and key label {DES:keylabel}. Retrieval will continue to be limited by the access controls in effect on the entry that contains the userPassword attribute values.

The *keylabel* must refer to a valid single-length data-encrypting key, also called a data key, generated by KGUP of ICSF and stored in the CKDS. The *keylabel* maximum length is 64 characters. See the information on managing cryptographic keys and using the Key Generator Utility Program in *z/OS: ICSF Administrator's Guide* for instructions on how to generate, store into CKDS, and authorize a data key and key label for DES encryption. It is important to remember to refresh both CKDS and RACF after you enter and authorize a key.

Notes:

- 1. When an encrypted password is stored in the TDBM backend, it is prefixed with the appropriate encryption tag so that when a clear text password is sent on an LDAP API simple bind it can be encrypted in that same method for password verification.
- The crypt() algorithm, implemented across many platforms, accepts only the first eight characters of a password. As a result, any password supplied on an Idap_simple_bind operation or Idap_compare operation that matches the first eight characters of a userPassword attribute value encrypted with the crypt() algorithm in the directory will match.
- The values returned by the crypt() algorithm are not portable to other X/Open-conformant systems. This means that user password values encoded by the crypt() algorithm and unloaded as tagged output using tdbm2ldif -t are not portable when loaded by another platform's load utility.

If **pwEncryption** is not specified in the configuration file, no password encryption takes place. User passwords are stored in clear text format and no tag is prefixed.

readOnly {on | off}

Global	TDBM	SDBM	EXOP
	X		

Specifies the ability to modify the database. Any attempt to use the LDAP server to modify the database will fail if **readOnly** is turned **on**.

Default = off

referral Idap_URL

Global	TDBM	SDBM	EXOP
X			

Specifies the referral to pass back when the local database cannot handle the request. It is also known as the default referral. The **referral** option can appear multiple times and should list equivalent servers.

In the following example, myldap.server.com is the host name and 3389 is the port number of the LDAP URL:

referral ldap://myldap.server.com:3389

securePort int

Global	TDBM	SDBM	EXOP
X			

Note: The **securePort** option has been deprecated by the **listen** option. See page 59 for information on the **listen** option.

TCP/IP port used for SSL communications.

Default = 636 Range = 1 - 65535

It is advisable to reserve the port number chosen here in your TCP/IP profile data set. Consult *z/OS: Communications Server: IP Configuration Guide* for further information.

security {ssl | sslonly | none | nossl}

Global	TDBM	SDBM	EXOP
Х			

Note: The **security** option has been deprecated by the **listen** option. See page 59 for information on the **listen** option.

Specifies what type of communications will be accepted. The **ssl** setting indicates that the server will listen on the secure port as well as the non-secure port. The **sslonly** setting means that the server will listen only on the secure port. The **none** or **nossl** settings indicate that the server will listen only on the non-secure port.

Default = none

sendV3stringsoverV2as {UTF-8 | ISO8859-1}

Global	TDBM	SDBM	EXOP
X			

Specifies the output data format to use when sending UTF-8 information over the LDAP Version 2 protocol.

Default = UTF-8

See "UTF-8 data over the LDAP Version 2 protocol" on page 298 for more detailed information on the use of this setting.

serverEtherAddr mac address

	Global	TDBM	SDBM	EXOP
	X			

Specifies the Media Access Control (MAC) address used for entry UUID generation. This value must be unique for all LDAP servers in your enterprise. You must specify the MAC address if multiple LDAP servers will run on a (hardware) system. This applies if your LDAP servers are on different LPARs and also if two LDAP servers are on the same LPAR. You do not need to specify this field if this is the only LDAP server that will run on this (hardware) system.

The suggested form of the mac_address you place here is:

4xmmmmssssss

Where:

Х

Is a one-character LDAP number. If more than one LDAP server is operating on a CPU, specify a different x value for each server. If more than 16 LDAP servers are desired, then use a serial number and model number from a CPU that is not running an LDAP server. If another CPU is not available, then set the x, mmmm, and ssssss values from the MAC address on an old Ethernet card that is no longer being used or not used to run an LDAP server.

mmmm Is the four-digit model number of the CPU.

ssssss Is the six-digit serial number of the CPU.

Note that the allowable values for *x*, *mmmm*, and *ssssss* are: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, and F.

Following is an example using the serverEtherAddr option:

serverEtherAddr 4A123401234D

serverKrbPrinc kerberosldentity

1

Global	TDBM	SDBM	EXOP
X			

Specifies the server's Kerberos identity that was created in "Defining the Kerberos identity" on page 36. This value is used to acquire the server credentials. Following is an example: serverKrbPrinc LDAP/myhost.myrealm.com@realm.com

Note that the LDAP portion of the *kerberosldentity* must be uppercase. Also, the *hostname* portion of the server's kerberos principal name is the primary hostname of the system on which the LDAP server is running in DNS.

servername string

Global	TDBM	SDBM	EXOP
	X		

Specifies the name of the DB2 server location that manages the tables for the LDAP Server. This value must match the name of one of the DATA SOURCE stanzas that must be specified in the ODBC initialization data set which is specified by the **dsnaoini** option in the configuration file. See DB2 for OS/390 and z/OS: ODBC Guide and Reference for a description of the DSNAOINI ODBC initialization data set contents. Using the example DSNAOINI file in Figure 2 on page 15 the value of *string* for **servername** would be L0C1.

sizeLimit int

Global	TDBM	SDBM	EXOP
X	X	X	

Specifies the maximum number of entries to return from a search operation. The maximum number can be modified on a specific search request as described below.

0 = no limit

Default = 500

Range = 0 - 2147483647

This option applies to all backends, unless specifically overridden in a backend definition. Specifying this prior to a **database** line in the configuration sets the option for all backends. Specifying it after a **database** line sets the option just for the backend defined by the **database** line.

Note that the following behavior is used when referring to the **sizeLimit** parameter. This parameter is also valid on an **Idapsearch** from the client.

- If a client has passed a limit, then the smaller value of the client value, and the value read from slapd.conf will be used.
- If the client has not passed a limit, and has bound as the adminDN, then the value specified in slapd.conf will be ignored.
- If the client has not passed a limit, and has not bound as the adminDN, then the limit will be that which was read from the **slapd.conf** file.

Note: When using the RACF access support of the z/OS LDAP server (SDBM), the number of entries returned will be subject to the limits of the SDBM implementation. The minimum of all of the restrictions and sizeLimit settings and specifications is the value used for any one search. See Chapter 18, "Accessing RACF information" on page 205 for a description of the limits associated with SDBM.

sslAuth {serverAuth | serverClientAuth}

Global	TDBM	SDBM	EXOP
Х			

Specifies the SSL/TLS authentication method. The serverAuth method allows the LDAP client to validate the LDAP server on the initial contact between the client and the server. The serverAuth method is the default.

The serverClientAuth method allows the LDAP client to validate the LDAP server and the LDAP server to validate the LDAP client if the client sends its digital certificate on the initial contact between the client and the server.

Note: To allow clients to SASL bind to the LDAP server, it is necessary to configure the server with sslAuth serverClientAuth.

See "Setting up for SSL/TLS" on page 43 for more SSL/TLS information.

sslCertificate {certificateLabel | none}

Global	TDBM	SDBM	EXOP
X			

Specifies the label of the certificate that will be used for server authentication. The certificate is stored in the key database file which is created and managed using the **gskkyman** tool. See z/OS: System Secure Sockets Layer Programming for details on using the gskkyman tool. Default = none

The default certificate, marked in the key database file managed by the gskkyman tool, will be used for server authentication.

sslCipherSpecs {string | ANY}

Global	TDBM	SDBM	EXOP
X			

Specifies the SSL/TLS cipher specifications that will be accepted from clients. The cipher specification is a blank delimitted string that represents an ORed bitmask indicating the SSL/TLS cipher specifications that will be accepted from clients. Clients that support any of the specified cipher specifications will be able to establish an SSL/TLS connection with the server. Table 10 on page 69

page 69 shows a list of the cipher spec mask values and the related decimal, hexidecimal, and keyword values. Refer to *z/OS*: System Secure Sockets Layer Programming for a description of supported cipher specifications.

The cipher specification may be specified as follows:

- A decimal value (for example, 256)
- · A hexadecimal value (for example, x100)
- A keyword (for example, TRIPLE_DES_SHA_US)
- A construct of those values using plus and minus signs to indicate inclusion or exclusion of a value. For example,
 - 256+512 is the same as specifying 768, or x100+x200, or TRIPLE DES SHA US+DES SHA US
 - 2816 is the same as specifying ALL-RC2_MD5_EXPORT-RC4_MD5_EXPORT

Table 10. Supported ciphers

Cipher	Hexadecimal value	Decimal value
TRIPLE_DES_SHA_US	x0100	256
DES_SHA_EXPORT	x0200	512
RC4_SHA_US	x0400	1024
RC4_MD5_US	x0800	2048
RC2_MD5_EXPORT	x1000	4096
RC4_MD5_EXPORT	x2000	8192
RSA_AES_128_SHA	x4000	16384
RSA_AES_256_SHA	x8000	32768
ANY	xFF00	65280
ALL	xFF00	65280

Depending upon the level of System SSL support installed, some ciphers may not be supported. System SSL will ignore the unsupported ciphers. You should consult the System SSL documentation to determine the specific ciphers that your installation supports.

See "Setting up for SSL/TLS" on page 43 for more SSL\TLS information.

Default = ANY

sslKeyRingFile name

Global	TDBM	SDBM	EXOP
X			

Specifies either the path and file name of the SSL/TLS key database file for the server or the name of the RACF key ring for the server.

The file name must match the key database file name that was created using the **gskkyman** utility (see *z/OS: System Secure Sockets Layer Programming*). Also, see "Setting up for SSL/TLS" on page 43 for more SSL information.

The LDAP server supports the use of a RACF key ring. Specify the RACF key ring name for the sslKeyRingFile and comment out the sslKeyRingFilePW and sslKeyRingPWStashFile options to use this support. Also, see "Creating and using a key database or key ring" on page 45 for more information on using RACF key rings.

sslKeyRingFilePW string

Global	TDBM	SDBM	EXOP
X			

Specifies the password protecting access to the SSL/TLS key database file. The password string must match the password to the key database file that was created using the gskkyman utility (see z/OS: System Secure Sockets Layer Programming). Also, see "Setting up for SSL/TLS" on page 43 for more SSL information.

Note: Use of the sslKeyRingFilePW configuration option is strongly discouraged. As an alternative, use either the RACF key ring support or the sslKeyRingPWStashFile configuration option. This will eliminate this password from the configuration file.

Comment out the sslKeyRingFilePW and sslKeyRingPWStashFile options if you are using RACF for the key ring.

sslKeyRingPWStashFile filename

Global	TDBM	SDBM	EXOP
X			

Specifies an HFS file name where the password for the server's key database file is stashed. Use the full path name of the stash file in HFS for filename.

If this option is present, then the password from this stash file overrides the sslKeyRingFilePW configuration option, if present. Use the gskkyman utility with the -s option to create a key database password stash file. See "Setting up for SSL/TLS" on page 43 for more SSL information.

Comment out the sslKeyRingFilePW and sslKeyRingPWStashFile options if you are using RACF for the key ring.

suffix dn suffix

Global	TDBM	SDBM	EXOP
	Х	X	

Denotes the root of a subtree in the namespace managed by this server within this backend. This option may be specified more than once to indicate all the roots of the subtrees within this backend.

Overlapping suffixes are not permitted in the LDAP server. Two suffixes overlap if they are identical or if one suffix is an ancestor of the other suffix. For example:

```
suffix "ou=Server Group,o=IBM"
suffix "ou=Server Group,o=IBM"
```

would not be permitted because the suffixes are identical. The following suffixes will not be permitted because one suffix is a descendant of the other:

```
suffix ou=Server Group, o=IBM
suffix o=IBM
```

If the suffix contains a special character (according to RFC 2253), it must be prefixed by two backslashes (\\). Note that the double backslashes are only needed in the configuration file; in all other usages, the special character is usually prefixed by a single backslash. For example, to use the suffix o=MyCompany#1, you must specify o=MyCompany\\#1 in the configuration file and o=MyCompany\#1 in other places where the suffix is used (for example, in LDAP search requests). See Chapter 18, "Accessing RACF information" on page 205 for exceptions to this when using SDBM. See IETF RFC 2253 Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names for valid DN formats.

Domain Component naming as specified by RFC 2247 is also supported in the LDAP server. For example, the domain name ibm.com could be specified as the following suffix in the LDAP configuration file:

suffix "dc=ibm,dc=com"

This option applies only to the backend in which it is defined. The EXOP backend does not require a suffix option in the configuration file. If one is specified, it will be ignored.

When using SDBM, do not specify more than one suffix. An SDBM suffix should not contain an alias name for an attribute. For example, an SDBM suffix cannot use the **surName** attribute (it can use the sn attribute instead). Also, you can use a case-sensitive attribute in the suffix, but SDBM ignores case when processing the suffix.

supportKrb5 {yes | no}

Global	TDBM	SDBM	EXOP
Х			

Specifies if the LDAP server participates in Kerberos GSS API Authentication. If it participates, then Kerberos GSS API binds are accepted and information is stored in the server's root DSE. Default = **no**

sysplexGroupName *group_name*

Global	TDBM	SDBM	EXOP
Х			

Specifies the name of the application group within which this server instance becomes a member for purposes of dynamic workload balancing. All concurrently-running LDAP servers which participate in dynamic workload balancing within the same Parallel Sysplex (see page 75 for the description of a Parallel Sysplex) must use the same group name in their server configuration file. This name may be up to 18 characters in length, and must be unique among all application groups using Workload Manager services.

The sysplexGroupName and sysplexServerName keywords are corequisites, and Sysplex Workload Management features will only be enabled if both are present with non-null arguments. When Sysplex Workload Management features are enabled, the server is automatically assumed to be in multi-server mode, and replication will be disabled in this server.

If sysplexGroupName, sysplexServerName, and multiserver keywords are all omitted from the server configuration file, the server will operate in single-server mode and replication will be possible.

sysplexServerName *server_name*

Global	TDBM	SDBM	EXOP
Х			

Specifies the name of the server within the group_name sysplex group which participates in dynamic workload balancing. (See page 76 for the description of a sysplex.) All concurrently-running LDAP servers which participate in dynamic workload balancing within the same Parallel Sysplex are identified by server_names within a given group_name used in the sysplex. This name may be up to 8 characters in length, and must be unique within a given group name using Workload Manager services.

The sysplexGroupName and sysplexServerName keywords are corequisites, and Sysplex Workload Management features will only be enabled if both are present with non-null arguments. When Sysplex Workload Management features are enabled, the server is automatically assumed to be in multi-server mode, and replication will be disabled in this server.

If sysplexGroupName, sysplexServerName, and multiserver keywords are all omitted from the server configuration file, the server will operate in single-server mode and replication will be possible.

timeLimit int

Global	TDBM	SDBM	EXOP
Х	X	X	

Specifies the maximum number of seconds (in real time) the LDAP server will spend answering a search request or a Modify DN request. This maximum number can be modified on a specific search request as described below. If a request cannot be processed within this time, a result indicating an exceeded time limit is returned.

0 = no limitDefault = 3600Range = 0 - 2147483647

This option applies to all backends, unless specifically overridden in a backend definition. Specifying this prior to a database line in the configuration sets the option for all backends. Specifying it after a database line sets the option just for the backend defined by the database line.

Note that the following behavior is used when referring to the **timeLimit** parameter. This parameter is also valid on an Idapsearch from the client.

- If a client has passed a limit, then the smaller value of the client value, and the value read from slapd.conf will be used.
- If the client has not passed a limit, and has bound as the adminDN, then the value specified in **slapd.conf** will be ignored.
- · If the client has not passed a limit, and has not bound as the adminDN, then the limit will be that which was read from the **slapd.conf** file.

useNativeAuth {selected | all | off}

Global	TDBM	SDBM	EXOP
	X		

Enables native authentication in the TDBM backend. If the value is:

- selected, only entries with the ibm-nativeld attribute that are within the native subtrees (see nativeAuthSubtree option on page 63) use native authentication.
- all, all entries within native subtrees use native authentication. These entries can contain the **ibm-nativeld** or **uid** attribute to specify the RACF ID.

• off, no entries participate in native authentication.

Default = off

validateincomingV2strings {on | yes | off | no}

Global	TDBM	SDBM	EXOP
X			

Specifies whether the incoming strings are validated. If set to on, this setting limits the format of incoming string data sent over the LDAP Version 2 protocol to the IA5 character set (X'00'-X'7F' or "7-bit ASCII"). With this setting, textual data received on operations outside of the IA5 character set causes the operations to fail with LDAP_PROTOCOL_ERROR.

Default = on

Note that while supported, it is not recommended to run with this data filtering disabled.

verifySchema {on | off}

Global	TDBM	SDBM	EXOP
X			

Note: The **verifySchema** option is now ignored by the LDAP server.

waitingThreads int

Global	TDBM	SDBM	EXOP
X			

Note: The waitingThreads option is deprecated and ignored by the LDAP server. The commThreads (page 55) and idleConnectionTimeout (page 57) options have replaced the waitingThreads option.

Deprecated options

The listen parameter deprecates the security, port, and securePort options in the configuration file. If a listen option is specified in the configuration file with either security, port, or securePort, the listen will take precedence over what has been specified for the deprecated security, port, and securePort options. If using an earlier version of the configuration file which contains the security, port, or securePort options, the LDAP server will be configured to listen on the port numbers specified for securePort, port, or both, depending upon the security setting. However, it is highly recommended that the LDAP server be configured using the listen option. See the description of the listen option on page 59 and "Migrating from port, securePort, and security to listen option" on page 106 for more information.

lanored options

The replKeyRingFile and replKeyRingPW options are no longer necessary or evaluated by the LDAP server. These options should be removed from the configuration file. Use the sslKeyRingFile option to specify the key database file and use the sslKeyRingPWStashFile configuration option or the RACF key ring support for the password.

The maxThreads and waitingThreads options are no longer necessary or evaluated by the LDAP server. These options should also be removed from the configuration file. Use the commThreads option to set the number of threads initialized at server start-up for communicating with the clients. See the description of the commThreads option on page 55 and "Migrating from maxThreads and waitingThreads to commThreads option" on page 107 for more information.

The attribute, objectclass, and verifySchema options are no longer necessary or evaluated by the LDAP server. These options should be removed from the configuration file.

With the removal of the RDBM backend, rdbm can no longer be specified as the dbtype on the database option. The index, tbspaceentry, tbspacemutex, tbspace32k, and tbspace4k options are no longer necessary or evaluated. These options should be removed from the configuration file.

Configuration considerations

The following table shows all of the different options you have and the decisions you must make for your LDAP server configuration. It also shows where you can find the associated reference information to help you make these decisions.

Table 11. Configuration considerations

Dependency	More information
HFS environment versus PDS (dataset) environment Depending on whether you use an HFS or PDS environment, there are some areas to consider for each.	"Setting up and running the LDAP server in the z/OS shell" on page 97 and "Setting up and running the LDAP server as a started task" on page 93
Operational mode You need to determine the type of operational mode your LDAP server will run in. For example, single-server mode or multi-server mode with or without dynamic workload management enabled.	"Determining operational mode" on page 75
TDBM backend You can use a TDBM backend database based on DB2.	Chapter 16, "LDAP directory schema for TDBM" on page 159
SDBM backend You can use an SDBM backend database based on RACF.	"Setting up for SDBM" on page 42
You can use an EXOP backend to retrieve Policy Director data.	Chapter 22, "Using extended operations to access Policy Director data" on page 247
SSL/TLS If you want to protect LDAP access with Secure Socket Layer (SSL) or Transport Layer Security (TLS), your LDAP server can be configured to provide server and, optionally, client authentication.	"Setting up for SSL/TLS" on page 43
Password encryption Your LDAP server can prevent unauthorized access to user passwords in a TDBM backend database.	"Configuring for user password encryption" on page 47
Kerberos authentication You can enable GSS API Kerberos binds and configure identity mapping.	Chapter 19, "Kerberos authentication" on page 225
Native authentication You can enable and configure your directory to perform authentication using the Security Server.	Chapter 20, "Native authentication" on page 233
CRAM-MD5 and DIGEST-MD5 authentication The LDAP server can be configured to perform CRAM-MD5 and DIGEST-MD5 authentication binds.	Chapter 21, "CRAM-MD5 and DIGEST-MD5 Authentication" on page 243

Table 11. Configuration considerations (continued)

Dependency	More information	
Administrator DN and password	"Establishing the administrator	
Determine how to set up your administrator DN and password.	DN and password" on page 80	
Replication	See Chapter 24, "Replication" on	
To keep multiple databases in sync, you can use replication.	page 269 for more information.	
Referrals	See Chapter 25, "Referrals" on	
To refer clients to additional directory servers, use referrals.	page 277 for more information.	
LDAP entry UUID	See Chapter 12, "Running and	
To generate Media Access Control (MAC) address used for entry UUID	using the LDAP entry UUID utility" on page 141	

[&]quot;Example configuration scenarios" on page 82 has a variety of examples showing different LDAP server configurations.

Determining operational mode

Note: If you already have the LDAP server installed, see Chapter 10, "Migrating to a z/OS LDAP server" on page 103.

Once the software has been installed, you are ready to configure it for use at your site. The LDAP server may be configured to run in one of several operational modes when a TDBM backend is configured.

Single-server mode

In this operational mode, only a single instance of the LDAP server may use a given TDBM database to store directory data. This server may perform replication (see Chapter 24, "Replication" on page 269) of database changes to other servers (on the same host system or on another host system).

See "Operating in single-server mode" on page 76 for more information.

Multiple single-server mode LDAP servers

In this operational mode, two or more LDAP servers, each in single-server mode, can be run on the same system with different TDBM backends. These servers may perform replication (see Chapter 24, "Replication" on page 269) of database changes to other servers (on the same host system or on another host system). However, each server must have its own separate replica.

See "Setting up multiple LDAP servers" on page 79 for more information.

Multi-server mode without dynamic workload management enabled

In this operational mode, multiple concurrent instances of the LDAP server using the same database to store directory data may run on a given host system, as well as on different host systems when those hosts are coupled in a Parallel Sysplex. A Parallel Sysplex is a collection of z/OS systems that cooperate, using certain hardware and software products, to process work. A Parallel Sysplex enables high-performance, multisystem data sharing across multiple Central Processor Complexes and z/OS images, as well as dynamic workload balancing across constituent systems in the sysplex. For additional information, see z/OS: Parallel Sysplex Overview: An Introduction to Data Sharing and Parallelism.

Multi-server mode is intended for use in an environment where high transactional volume is common, or where maximum availability is required. This mode provides benefits of improved availability, fault tolerance, improved resource utilization, and improved performance. These benefits are achieved by enabling concurrent running of multiple servers which are functionally equivalent and which provide access to the same LDAP data.

Note: In this operational mode, replication of TDBM database changes to other servers is not supported.

See "Operating in multi-server mode without dynamic workload management enabled" on page 77 for more information.

Multi-server mode with dynamic workload management enabled

This operational mode augments the benefits of the previously described mode with dynamic workload management. This mode may only be used when all host systems on which instances of the LDAP server will run are coupled in the same Parallel Sysplex and all instances of the LDAP server are using the same TDBM database. The dynamic workload management for LDAP servers is provided through the use of an z/OS TCP/IP feature called "connection optimization". Connection optimization uses Domain Name Services (DNS) for distributing connections among server applications within a sysplex domain. Connection optimization achieves workload balancing by distributing connections to systems with the most available resources and by avoiding unavailable sysplex resources. See z/OS: Communications Server: IP Configuration Reference for information on connection optimization.

Note: In this operational mode, replication of TDBM database changes to other servers is not supported.

See "Operating in multi-server mode with dynamic workload management enabled" on page 77 for more information.

· Program call (PC) callable support mode

The program call (PC) callable support in LDAP provides a program call interface to the LDAP extended operations backend (EXOP). This interface is only available using the z/OS SAF interfaces designed to allow Policy Director access to LDAP data.

See "Operating in PC callable support mode" on page 79 for more information.

In any of these modes, all combinations of TDBM (one or more), SDBM, and EXOP backends are supported.

Notes:

- 1. A single LDAP server instance can have one SDBM backend, but it can have multiple TDBM backend instances
- 2. If multiple single-server mode LDAP servers are being used on the same system, only one of the LDAP servers can be configured for PC callable support.
- 3. If either multi-server mode is being used and RACF data will be accessed from both servers, then the RACF database should also be shared across the systems where the LDAP servers run to ensure consistency of SDBM operations.

Operating in single-server mode

For the LDAP server to operate in single-server mode, the server configuration file may contain any of the previously documented options except the sysplexGroupName and sysplexServerName options (the presence of which causes the LDAP server to operate in multi-server mode with dynamic workload management enabled). If the multiserver option is present, its value must be set to either n or N.

Restrictions

If one LDAP server instance using a given DB2-based backend to store directory information is operating in single-server mode, it must be the *only* instance of the LDAP server using that DB2-based backend. Configuring more than one LDAP server instance to use the same DB2 database may yield unpredictable results if one or more of those server instances is configured in single-server mode. If it is desired to access the same DB2-based backend from more than one server instance, all server instances using the same DB2-based backend must be configured to operate in multi-server mode.

Operating in multi-server mode without dynamic workload management enabled

For the LDAP server to operate in multi-server mode without dynamic workload management enabled, the server configuration file may contain any of the previously documented options except the sysplexGroupName and sysplexServerName options. If either of these keywords are present, they cause the LDAP server to operate in multi-server mode with dynamic workload management enabled. The multiserver mode option must be present with a value of y or Y. All instances of the LDAP server using the same DB2-based backend must have the multiserver option present in the configuration file used to start each server instance when operating in this mode.

When you are using referrals without dynamic workload management enabled, multiple default referrals defined for other servers can be set up to point to each of the multiple server instances. Similarly, any referral objects defined in other servers which point to the multiple server instances can have multi-valued **ref** attributes set up, each of which is an LDAP URL pointing to the corresponding server instances.

Dependencies

The LDAP server may operate in multi-server mode without dynamic workload management enabled. It is still possible to run multiple concurrent server instances configured to use the same DB2-based backend on the same z/OS image, and/or run multiple concurrent server instances on multiple z/OS images coupled in a Parallel Sysplex. If multiple instances will be run on multiple z/OS images in a Parallel Sysplex, the DB2 subsystems to which each server instance will attach (see "Getting DB2 installed and set up for CLI and ODBC" on page 13) must be configured on each of the images as members of the same DB2 data sharing group. (See DB2 for OS/390 and z/OS: Data sharing: Planning and Administration for information on planning, installing, and enabling DB2 data sharing, and z/OS: MVS Setting Up a Sysplex for information on planning and installing a Parallel Sysplex.)

Restrictions

If multiple LDAP server instances are using the same DB2-based backend to store directory information, all LDAP server instances using that database must be operating in multi-server mode, either with or without dynamic workload management enabled.

Note: When configuring multiple servers to access the same set of DB2 tables, the dbuserid option must be set to the same value for all the servers that are accessing the same tables.

LDAP server instances operating in multi-server mode, either with or without dynamic workload management enabled, will not perform replication (see Chapter 24, "Replication" on page 269), even if replication objects are present in the DB2-based backend (TDBM). If replication is required, single-server mode must be used.

Operating in multi-server mode with dynamic workload management enabled

For the LDAP server to operate in multi-server mode with dynamic workload management enabled, the server configuration file may contain any of the previously documented options. The sysplexGroupName and sysplexServerName options must both be present in the server configuration file. If the multiserver option is present with a value of n or N, the option value is overridden and treated as though it were set to

To exploit the dynamic workload management feature, the z/OS images on which the LDAP server runs must be coupled in a Parallel Sysplex. Name servers must be configured for connection optimization and started in the same sysplex in which the LDAP server instances are running. See "Dependencies" on page 78 for more information.

When multiple concurrent instances of the LDAP server are operating in multi-server mode with dynamic workload management enabled, server instances within the same group are considered to provide

equivalent service. In essence, the servers are treated as clones of each other. With this in mind, it should be noted that the port on which the server is started should be the same for each instance.

To connect to any unspecified server instance in sysplexGroupName group_name, the client specifies a target host name of group_name.sysplex_domain_name, where group_name is the name of the application group in which servers are configured using the sysplexGroupName option in the server configuration file used to start each respective LDAP server instance, and sysplex_domain_name is the name or alias for the sysplex domain. The client will be connected to a server instance in the group on the system in the sysplex with the most available resources, at the port specified by the client which must agree with the port configured when the server instances in the group were started. While it is possible to specify a particular server instance using a fully-qualified server name, doing so reduces the effectiveness of dynamic workload management through connection optimization.

Also note that it is possible to run multiple concurrent LDAP servers using the same DB2-based backend with a mix of servers enabled or not enabled for dynamic workload management, but doing so also reduces the effectiveness of dynamic workload management through connection optimization.

When you are using referrals with dynamic workload management enabled, a single default referral defined at other servers is used where the host is specified as sysplexGroupName.sysplexDomainName and similarly, referral objects defined in other servers use this as the host within a single-valued ref attribute.

Dependencies

The LDAP server may operate in multi-server mode with dynamic workload management enabled while running multiple concurrent server instances configured to use the same DB2-based backend on multiple z/OS images coupled in a parallel sysplex. The DB2 subsystems to which each server instance will attach (see "Getting DB2 installed and set up for CLI and ODBC" on page 13) must be configured on each of the images as members of the same DB2 data sharing group. (See DB2 for OS/390 and z/OS: Data sharing: Planning and Administration for information on planning, installing, and enabling DB2 data sharing, and z/OS: MVS Setting Up a Sysplex for information on planning and installing a Parallel Sysplex.)

Name servers must be configured for connection optimization and started in the same sysplex in which the LDAP server instances are running. Proper distribution of server application addresses for connection optimization to function properly requires DNS queries to be answered by the name server within the sysplex. For this reason, name servers located outside the sysplex cannot be configured as primary or secondary servers for the sysplex domain.

The port numbers on which the LDAP server instances will listen, both secure and unsecure, must be the same for all servers in the same sysplexGroupName for the dynamic workload management to be effective.

Also, the Workload Management Services on each host in the sysplex must be configured in "goal mode". (See z/OS: Communications Server: IP Configuration Reference for information on configuring a sysplex domain for connection optimization and for information on how to configure Workload Management Services in goal mode.)

Restrictions

If multiple LDAP server instances are using the same DB2-based backend to store directory information, all LDAP server instances using that database must be operating in multi-server mode, either with or without dynamic workload management enabled.

Note: When configuring multiple servers to access the same set of DB2 tables, the dbuserid option must be set to the same value for all the servers that are accessing the same tables.

LDAP server instances operating in multi-server mode, either with or without dynamic workload management enabled, will not perform replication (see Chapter 24, "Replication" on page 269), even if replication objects are present in the DB2 database.

Operating in PC callable support mode

The program call (PC) callable support in LDAP provides a program call interface to the LDAP extended operations backend (EXOP). This interface is only available using the z/OS SAF interfaces designed to allow Policy Director access to LDAP data. The PC callable support is initialized in an LDAP server when the appropriate listen record is included in the configuration file or specified when starting the server. An LDAP server can be dedicated to running just the PC callable support or it can run the PC callable support in addition to its normal socket interfaces.

Running the PC callable support has two interactions with the system:

- The address space of the LDAP server is made non-swappable during initialization of the PC callable support. As a result, resources used by that address space can significantly affect system performance.
- Because the PC callable support connects its PC table to a system index, the address space identifier of the LDAP server address space is not re-usable until the next IPL. If the system is configured with a low limit on the number of address spaces, it is possible to run out of address space identifiers, preventing new address spaces from being started. This problem is more likely to occur if the LDAP server running PC callable support is frequently brought down and re-started.

Consider configuring a separate LDAP server to run only the PC callable support. Because the server is not also running the backend controlling the data, fewer resources will be made non-swappable and the server will less likely need to be re-started. The disadvantage is that an extended operation request will require the LDAP server to communicate with another LDAP server for the data needed to satisfy the request, which can be slower than accessing that data on the same LDAP server. In general,

- If the data used in the extended operations is not on this system, then configure a separate LDAP server for the PC callable support.
- If the data is on this system, then try both configurations (PC callable support in a separate LDAP server and PC callable support in the same LDAP server as the data) to determine the impact on performance.

At most one LDAP server in a system can activate PC callable support. If an LDAP server tries to initialize PC callable support after another LDAP server has already tried (successfully or unsuccessfully) to initialize PC callable support, the initialization fails. The first LDAP server that tries to initialize the PC callable support locks the access to the PC callable support until that LDAP server has been shut down. If you are running LDAP in a sysplex, configure one LDAP server on each system in the sysplex to run the PC callable support. Each system should share the DB2 and RACF databases to ensure that they return the same results.

Setting up multiple LDAP servers

In order to set up two or more LDAP servers on the same system with different DB2-based backends, do the following:

- Follow the steps outlined in "Creating the DB2 database and table spaces for TDBM" on page 40 and be sure to:
 - 1. Modify the SPUFI file that creates tables and indexes. Change -UUUUUUUU- (database owner) and -DDDDDDDD- (database name) to a different value and submit the SPUFI. A separate set of DB2 tables will be created.
 - 2. Update the configuration file **dbuserid** option with the value from step 1

Establishing the administrator DN and password

There are several ways that the administrator DN and password or the master server DN and password can be configured. One of these ways must be used, since an administrator DN and password are required for the LDAP server and some other LDAP programs to operate. The administrator DN must be present in the configuration file using the adminDN option (see page 54). The administrator DN password can optionally (this is not recommended) be placed in the configuration file using the adminPW option (see page 54) or can be held in the namespace managed by this instance of the LDAP server. If a replica is being established, the masterServerDN option must be present in the configuration file. The masterServerPW option can optionally be present. (This is also not recommended.) All of the options described below are applicable for adminDN and the first two options described below are applicable for masterServerDN.

· Administrator DN and password in configuration file

The simplest but least secure method is to select an administrator DN that is outside of the scope of suffixes managed by this server (see the suffix option on page 70). In other words, choose an administrator DN such that it does not fall within the portion or portions of the namespace managed by this server. Selection of this type of administrator DN requires that the password be placed in the configuration file using the adminPW option (see page 54).

For example, you might choose a simple DN, such as "cn=Admin" for the administrator DN and a simple password such as secret. The configuration file options would then be established this way:

```
adminDN "cn=Admin"
adminPW secret
```

Note: Do not use the example above without changing the password value, as well as the actual distinguished name.

When a program or user binds using this administrator DN, the LDAP server verifies that the password supplied on the request matches the value provided in the configuration file for the adminPW option.

Note: When first configuring a TDBM backend, it may be necessary to use this approach until the schema supporting the directory entries is loaded. Once the schema is loaded and the entry representing the administrator is added, the adminDN can be changed to the entry DN (see the next list item regarding "Administrator DN and password as a TDBM entry"). The server must be restarted to pick up the new adminDN. Alternately, use the Idif2tdbm utility to load both the schema and adminDN entry.

Administrator DN and password as a TDBM entry

In this method, the administrator DN is established as an entry managed by the TDBM backend. The userPassword attribute is used to hold the password for the administrator DN in this case.

For example, if the TDBM database is managing the portion of the namespace "o=Your Company", one administrator DN that could be selected would be "cn=LDAP Admin,o=Your Company".

The configuration file would include the following options:

```
adminDN "cn=LDAP Admin,o=Your Company"
database tdbm GLDBTDBM
suffix "o=Your Company"
```

The LDIF-format entry to be added to the database through Idapadd or Idif2tdbm might be:

```
dn: cn=LDAP Admin,o=Your Company
objectclass: person
cn: LDAP Admin
description: Administrator DN for o=Your Company server
sn: Administrator
uid: admin
userpassword: secret
```

Note: Do not use the example above without changing the password value, as well as the actual distinguished name.

If this entry is contained in an HFS file called admin.ldif, it can be loaded using **Idapadd**: ldapadd -h *Idaphost* -p *Idapport* -D *binddn* -w *passwd* -f admin.ldif

Note that the LDIF-format entry can alternately be stored in a dataset and loaded through TSO using the sample CLIST for **Idapadd**.

Note: The Idapadd example above assumes that the LDAP server is running and that the "suffix entry" (entry with the name "o=Your Company") already exists. Furthermore, the binddn is assumed to exist and have sufficient authority to add the entry. When initially setting up the LDAP server, one way to satisfy the assumption is to first configure the LDAP server using the adminDN and adminPW configuration options. Then start the LDAP server, load the "suffix entry" and the "admin entry", using binddn and passwd as the adminDN and adminPW configuration values, respectively. After the add operations complete, stop the LDAP server, change the adminDN configuration option value to the name of the entry just added and remove the adminPW configuration option. Then restart the LDAP server.

As an alternative to the steps listed in the previous paragraph, you can use the load utility (**Idif2tdbm** for TDBM) to load the admin DN entry.

When a program or user binds using this administrator DN, the LDAP server verifies that the password supplied on the request matches the value of the **userPassword** attribute stored in the entry in DB2.

CRAM-MD5 and DIGEST-MD5 authentication binds with the **adminDN** are supported as long as the entry exists in a TDBM backend. The **adminDN** entry must contain a **uid** attribute value that will be used as the user name by a client application when attempting a CRAM-MD5 or DIGEST-MD5 authentication bind. For more information on CRAM-MD5 and DIGEST-MD5 authentication, see Chapter 21, "CRAM-MD5 and DIGEST-MD5 Authentication" on page 243.

Administrator DN and password in RACF

This method requires that the LDAP server be configured to use the RACF support provided in the SDBM backend. The administrator DN can be established as a RACF-style DN based upon a RACF user ID. (See "RACF-style distinguished names" on page 156 for more information.) In this case, the password for the administrator DN is the RACF user ID's password, and is stored and verified by RACF. You cannot use the **krbLDAPAdmin** option to provide a Kerberos identity for the **masterServerDN** option.

For example, if you configure the LDAP server with RACF support where the portion of the namespace held by RACF is "sysplex=Sysplex1,o=Your Company", and the RACF user ID that is used for the administrator is gladmin, the configuration file would include these options:

```
adminDN "racfid=gladmin,profiletype=user,sysplex=Sysplex1,o=Your Company"
...
database sdbm GLDBSDBM
suffix "sysplex=Sysplex1,o=Your Company"
```

When a program or user binds using this administrator DN, the LDAP server makes a request to RACF to verify that the password supplied on the request matches the RACF password for RACF user ID gladmin.

Note that the user ID specified must have an OMVS segment defined and an OMVS UID present.

• **krbLDAPAdmin** in the configuration file

You may wish to configure the administrator to be able to bind to the server through Kerberos. In this case you need to create a Kerberos identity for the user and also add this value to the configuration file. For example, if the user ID for the LDAP administrator is **LDAPADM** and this Kerberos identity was configured to be <code>ldapadm@realm1.com</code> then your configuration file will have the following:

This allows the administrator to bind to the server through Kerberos Authentication rather than by performing a simple bind.

Example configuration scenarios

This section shows scenarios of LDAP server configurations.

Configuring a TDBM backend with SSL/TLS only and password encryption

The configuration example in this section uses the TDBM backend and shows the configuration file checklist next to the corresponding sample configuration file.

Table 9 on page 52 shows all the options that are available for each section.

Table 12. Sample checklist and sland conf (using TDBM, SSI /TLS, and password encryption)

Section	Check	Sample slapd.conf	
Global	√	# Filename slapd.conf	
SSL/TLS	V	# Global section	
Sysplex		sizelimit 500 timelimit 3600	
Kerberos		adminDn "cn=LDAP Administrator,o=Your Company"	
SDBM backend		listen ldaps://:636	
Kerberos		sslAuth serverClientAuth sslCertificate none sslCipherSpecs 15104 sslKeyRingFile /u01/ldapsrv/ldapsrv.kdb sslKeyRingPWStashFile /u01/ldapsrv/ldapsrv.sth	
TDBM backend	V		
Password encryption	V		
Replication		# TDBM backend section database tdbm GLDBTDBM LocalDirectory	
Multi-server		suffix "o=Your Company"	
Kerberos		servername LOC1 dbuserid LDAPSRV	
Native authentication		databasename LDAPDB attrOverflowSize 500	
EXOP backend		pwEncryption MD5	

Configuring SDBM and TDBM backends

The configuration example in this section uses both SDBM and TDBM backends and shows the configuration file checklist next to the corresponding sample configuration file.

Table 9 on page 52 shows all of the options that are available for each section.

Table 13. Sample checklist and slapd.conf (using SDBM and TDBM)

Section	Check	Sample slapd.conf
Global	V	# Filename slapd.conf
SSL/TLS		# Global section
Sysplex		sizelimit 500
Kerberos		timelimit 3600 adminDn "racfid=ldadmin,profiletype=user,cn=myRACF"
SDBM backend	V	
Kerberos		# SDBM backend section database sdbm GLDBSDBM
TDBM backend	V	suffix "cn=myRACF"
Password encryption		# TDBM backend section
Replication		database tdbm GLDBTDBMsuffix "o=Your Company"
Multi-server		servername LOC1
Kerberos		dbuserid LDAPSRV databasename LDAPDB
Native authentication		attrOverflowSize 500
EXOP backend		

Configuring an EXOP backend

The configuration example in this section uses an EXOP backend and shows the configuration file checklist next to the corresponding sample configuration file.

Table 9 on page 52 shows all of the options that are available for each section.

Table 14. Sample checklist and slapd.conf (using EXOP)

Section	Check	Sample slapd.conf
Global	V	# Filename slapd.conf
SSL/TLS		# Global section
Sysplex		listen ldap://:pc
Kerberos		# EXOP backend section
SDBM backend		database exop GLDXPDIR
Kerberos		
TDBM backend		
Password encryption		
Replication		
Multi-server		
Kerberos		
Native authentication		
EXOP backend	V	

Configuring and using multiple concurrent servers in a sysplex

Following is a simplified example scenario to demonstrate the configuration and usage of the multi-server operation mode with dynamic workload balancing enabled. While the user's operational environment may be more complex than this example, the lessons demonstrated apply in a similar fashion.

See "Determining operational mode" on page 75 for more information about operating in a sysplex.

Scenario

ABC Company is running a 3-way Parallel Sysplex with DB2 data sharing available on each host in the sysplex. Configure an instance of the LDAP server on each of the z/OS systems in the sysplex, serving the same LDAP directory TDBM database, such that the three instances are functional equivalents of each other, permitting users to exploit dynamic workload balancing across these LDAP servers in the sysplex.

Assume that at this point, the system administrator has installed and configured a DB2 subsystem on each host in the sysplex to be part of the same data sharing group, and that the system administrator has configured at least one Domain Name Service (DNS) server for the sysplex. In addition, the TCP/IP stacks which serve each host are registered with Workload Manager (WLM). See "Dependencies" on page 78 for more information.

The operational environment in which the LDAP servers will run is configured as indicated in the following diagram:

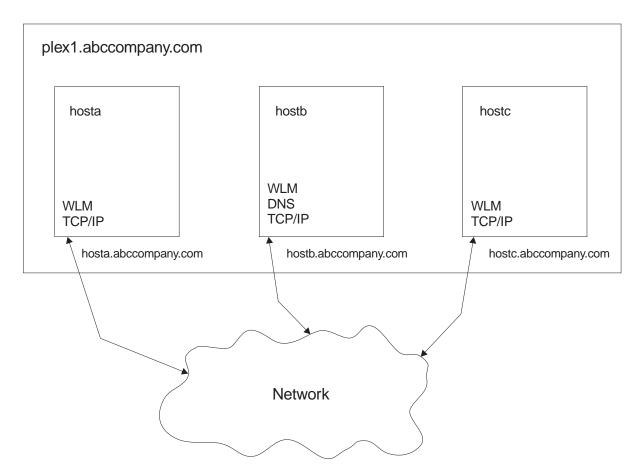


Figure 7. Multi-server sample configuration (phase 1)

The three host systems are named hosta, hostb, and hostc, and are coupled in sysplex plex1 in Internet sub-domain abccompany.com. Each of the host systems is configured with a single network adapter, and these systems are accessible with these names:

hosta.abccompany.com hostb.abccompany.com hostc.abccompany.com

The sysplex domain name for this sysplex is plex1.abccompany.com. The primary DNS server for the sysplex domain runs on hostb.

Three server instances will be started, one on each host in the sysplex. A checklist and the server configuration file for each of the three servers follows and differences among the files are highlighted.

Table 15. Sample checklist (using TDBM with multi-server and dynamic workload management)

Section	Check
Global	V
SSL/TLS	
Sysplex	V
Kerberos	
SDBM backend	
Kerberos	
TDBM backend	V
Password encryption	
Replication	
Multi-server	V
Kerberos	
Native authentication	

```
# * Filename slapd.conf
listen 1dap://:389
listen ldaps://:636
commThreads
                60
                40000
maxconnections
timelimit
                3600
sizelimit
                500
sysplexGroupName ldapgrp1
sysplexServerName servera
# tdbm database definitions
tdbm GLDBTDBM
database
suffix
            "o=Your Company"
# The following options must be filled in with appropriate values
# for your DB2 setup, prior to attempting to run with the DB2 backend.
multiserver
servername
            loc1
databasename abcdb1
dbuserid
            dbu01
dsnaoini ABCCO.DB2CLI.CLIINIA
readOnly
            off
```

Figure 8. Configuration file for Server A on hosta

Figure 9 shows the contents of ABCCO.DB2CLI.CLIINIA:

```
;This is a comment line...
; Example COMMON stanza
[COMMON]
MVSDEFAULTSSID=DB1G
; Example SUBSYSTEM stanza for your DB2 subsystem name
[DB1G]
MVSATTACHTYPE=RRSAF
PLANNAME=DSNACLI
; Example DATA SOURCE stanza for your data source
[LOC1]
AUTOCOMMIT=0
CONNECTTYPE=1
```

Figure 9. Contents of ABCCO.DB2CLI.CLIINIA

```
# * Filename slapd.conf
listen ldap://:389
listen ldaps://:636
              60
commThreads
maxconnections
              40000
             3600
timelimit
sizelimit
              500
sysplexGroupName ldapgrp1
sysplexServerName serverb
# tdbm database definitions
database
          tdbm GLDBTDBM
suffix
          "o=Your Company"
# The following options must be filled in with appropriate values
# for your DB2 setup, prior to attempting to run with the DB2 backend.
multiserver
servername
              loc1
databasename
              abcdb1
              dbu01
dbuserid
dsnaoini ABCCO.DB2CLI.CLIINIB
readOnly
              off
```

Figure 10. Configuration file for Server B on hostb

Figure 11 on page 88 shows the contents of ABCCO.DB2CLI.CLIINIB:

```
;This is a comment line...
; Example COMMON stanza
[COMMON]
MVSDEFAULTSSID=DB2G
; Example SUBSYSTEM stanza for your DB2 subsystem name
MVSATTACHTYPE=RRSAF
PLANNAME=DSNACLI
; Example DATA SOURCE stanza for your data source
[LOC1]
AUTOCOMMIT=0
CONNECTTYPE=1
```

Figure 11. Contents of ABCCO.DB2CLI.CLIINIB

```
# * Filename slapd.conf
listen ldap://:389
listen ldaps://:636
commThreads
maxconnections
              40000
             3600
timelimit
sizelimit
              500
sysplexGroupName ldapgrp1
sysplexServerName serverc
# tdbm database definitions
database
            tdbm GLDBTDBM
suffix
            "o=Your Company"
# The following options must be filled in with appropriate values
# for your DB2 setup, prior to attempting to run with the DB2 backend.
multiserver y
servername
         loc1
databasename abcdb1
dbuserid dbu01
dsnaoini ABCCO.DB2CLI.CLIINIC
readOnly
          off
```

Figure 12. Configuration file for Server C on hostc

Figure 13 on page 89 shows the contents of ABCCO.DB2CLI.CLIINIC:

```
;This is a comment line...
; Example COMMON stanza
[COMMON]
MVSDEFAULTSSID=DB3G
; Example SUBSYSTEM stanza for your DB2 subsystem name
[DB3G]
MVSATTACHTYPE=RRSAF
PLANNAME=DSNACLI
; Example DATA SOURCE stanza for your data source
[LOC1]
AUTOCOMMIT=0
CONNECTTYPE=1
```

Figure 13. Contents of ABCCO.DB2CLI.CLIINIC

The DB2 subsystem IDs must be different on each system and those subsystem IDs must all be defined into the same DB2 data sharing group.

Several things should be noted in the server configuration files above:

- All three configuration files use the same TDBM database (which is accessible through a DB2 data sharing group which contains the subject database) and database resources, as indicated by identical values for the parameters **servername**, **databasename**, and **dbuserid**. Each server configuration file points to the DB2 Call Level Interface (CLI) initialization file for that respective server. The MVSDEFAULTSSID defined in the CLI initialization file for each server must be the name of the DB2 subsystem on the host on which that server instance will be started which is a member of the DB2 data sharing group which all three systems in the sysplex share, and which contains the TDBM database of interest.
- All three configuration files use the same name for their sysplexGroupName option; this is required to
 ensure all three servers are recognized as functional equivalents of each other for purposes of dynamic
 workload management. In addition, the sysplexServerName for each of the three must be unique
 within this sysplexGroupName in this sysplex.
- The ports (both secure and nonsecure) on the listen parameters in the configuration file must be the same for all servers in the same sysplexGroupName for the dynamic workload management to be effective.

With the additional configuration information just outlined, we can extend the diagram in Figure 7 on page 85 to look like this:

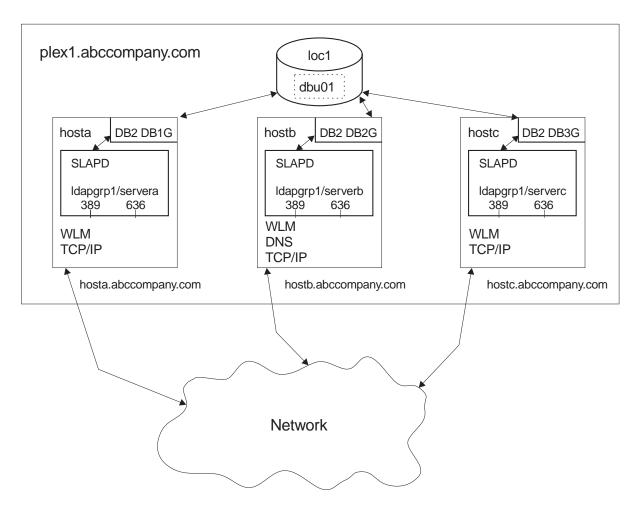


Figure 14. Multi-server sample configuration (phase 2)

When the LDAP servers are started, the following messages will be printed to STDERR:

GLD0115I Workload Manager enablement initialization successful for group=ldapgrp1, server=servera on host HOSTA.

GLD0115I Workload Manager enablement initialization successful for group=ldapgrp1, server=serverb on host HOSTB.

GLD0115I Workload Manager enablement initialization successful for group=ldapgrp1, server=serverc on host HOSTC.

At this point, a client can send a search request to any LDAP server instance which will provide service for the TDBM database of interest by using the sysplex group name in the DNS host name option as in:

```
ldapsearch -h ldapgrp1.plex1.abccompany.com -p 389 -D "cn=admin" -w secret
 -b "o=ABCCOMPANY" objectclass=person cn
```

With the host specified as ldapgrp1.plex1.abccompany.com, the request will be routed to the server instance in sysplexGroupName | dapgrp1 in sysplex plex1.abccompany.com which has the most available resources with which to perform the work (see Figure 14). Although requests could be directed at a specific server instance (that is, specifying serverc.ldapgrp1.plex1.abccompany.com or hostc.abccompany.com for the DNS host name), doing so defeats the use of the dynamic workload management features by bypassing the TCP/IP connection optimization.

See z/OS: Security Server LDAP Client Programming for more information about Idapsearch.

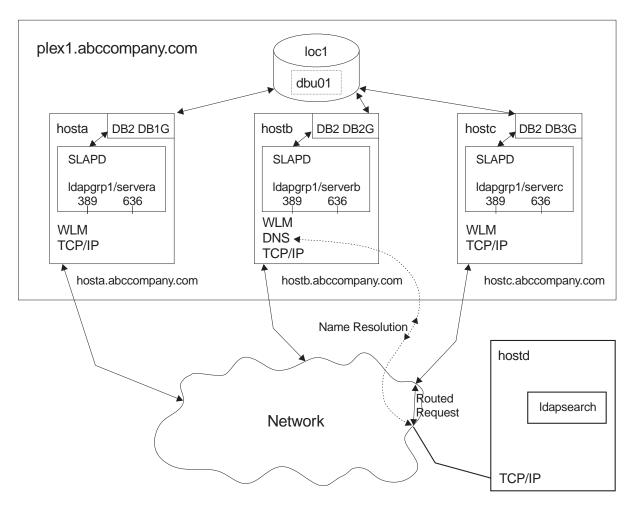


Figure 15. Multi-server sample configuration (phase 3)

It should be noted that for proper workload management using TCP/IP connection optimization, DNS queries must be answered by the name server within the sysplex. For this reason, name servers that are located outside the sysplex cannot be configured as primary or secondary servers for the sysplex domain. Thus, in this example, hostd must be configured to resolve names through the sysplex name server.

Chapter 9. Running the LDAP server

This chapter describes what is necessary to get the LDAP server running:

- Be sure your LDAP server can locate the DLLs
- · Set up and run the LDAP server as a started task
- Set up and run the LDAP server in the z/OS shell
- · Use dynamic debugging
- · Gather trace records in the in-storage trace table
- · Capture performance information
- · Verify the LDAP server

Note: The LDAP server relies on TCP/IP and DB2 (for TDBM) being available. If either of these products fail while the LDAP server is running, the LDAP server must be restarted.

Setting up the PDS for the LDAP server DLLs

The LDAP server searches for and loads a number of dynamic load libraries (DLLs) during its startup processing. All DLLs for the LDAP server are shipped in PDS format only. In order for these DLLs to be located by the LDAP server at runtime, the PDS which contains these DLLs (*GLDHLQ*.SGLDLNK) must either be in the LINKLIST, referenced in a **STEPLIB DD** card, if the LDAP server is started from JCL, or listed in the **STEPLIB** environment variable, if the LDAP server is started from the USS (Unix System Services) command prompt. Any of these methods can be used, and the choice of the best method is dependent on the way you will most often be running the LDAP server. Note that if you put *GLDHLQ*.SGLDLNK in LINKLIST, **STEPLIB** is not necessary.

The LDAP server also depends on the **SCEERUN** dataset. Add **SCEERUN** to your LINKLIST or, if that is not possible, add it to STEPLIB. See *z/OS: Program Directory* for more information.

Note: If you have a ServerPac installation, *GLDHLQ* will be **GLD**.

Setting up and running the LDAP server as a started task

To run the LDAP server as a started task, you must define the started task for the LDAP server and then you can run the LDAP server using JCL.

Defining the started task for the LDAP server

After you create the **LDAPSRV** user ID (described in "Requirements for a user ID that runs the LDAP server" on page 35), you must define the **LDAPSRV** started task. The examples and the sample startup procedure use the name **LDAPSRV** for this task, but you can use any name for it.

To define the started task for the user ID you just created, you can use the following RACF commands.

RDEFINE STARTED LDAPSRV.** STDATA(USER(LDAPSRV))

SETROPTS RACLIST(STARTED) REFRESH

Running the LDAP server using the sample JCL

The JCL needed to run the LDAP server as a started task is provided with the product as a procedure. This JCL can be found in *GLDHLQ*.SGLDSAMP on the system where the LDAP server is installed. (The sample JCL is shown in "Sample JCL for the LDAP server" on page 413.) This JCL procedure can be started in the System Display and Search Facility (SDSF) or from the operator's console, once the sample JCL has been placed into the installation-specific library for procedures. This JCL must be tailored before it can be run.

To start the LDAP server in SDSF, enter:

/s ldapsrv

To start the LDAP server from the operator's console, enter:

s ldapsrv

The LDAP server has the following optional command-line parameters. One or more of these may be specified when starting the LDAP server.

-f pathname

Name of configuration file to be read. Default is /etc/ldap/slapd.conf.

-I Idap URL

Host name or IP address and port number on which the LDAP server should bind and listen for incoming requests. Refer to the listen parameter on page 59 for information on the Idap_URL parameter.

It is advisable to reserve the port number or numbers chosen here in your TCP/IP profile data set. Also, be aware that port numbers below 1024 may require additional specification. See z/OS: Communications Server: IP Configuration Guide.

See "Migrating from slapd command-line arguments -p and -s to -l" on page 107 for migration information.

-p port

Note: The port option is deprecated when the listen option is specified. See page 59 for information on the **listen** option.

Port number where the LDAP server will listen for nonsecure communications. If you specify this parameter, it overrides the value that may be in the **slapd.conf** configuration file if there are not any listen parameters specified in the configuration file or on the command line.

-s secureport

Note: The port option is deprecated when the listen option is specified. See page 59 for information on the **listen** option.

Port number where the LDAP server will listen for secure communications. If you specify this parameter, it overrides the value that may be in the slapd.conf configuration file if there are not any listen parameters specified in the configuration file or on the command line.

-d debug level

The debug level is a mask that may be specified as follows:

- A decimal value (for example, 32)
- A hexadecimal value (for example, x20 or X20)
- A keyword (for example, FILTER)
- A construct of those values using plus and minus signs to indicate inclusion or exclusion of a value. For example:
 - '32768+8' is the same as specifying '32776', or 'x8000+x8', or 'ERROR+CONNS'
 - '2146959359' is the same as specifying 'ANY-STRBUF'
 - By beginning the debug level with a minus sign, you can deactivate debug collection for the various types. "-CONNS" modifies an existing debug level by deactivating connection traces.
 - By beginning the debug level with a plus sign, you can activate debug collection for the various types. "+CONNS" modifies an existing debug level by activating connection traces.

Table 16 on page 95 lists the debug levels and the related decimal, hexadecimal, and keyword values.

Table 16. Debug levels

Keyword	Decimal	Hexadecimal	Description		
OFF	0	0x00000000	No debugging		
TRACe	1	0x00000001	Entry and exit from routines		
PACKets	2	0x00000002	Packet activity		
ARGS	4	0x00000004	Data arguments from requests		
CONNs	8	0x00000008	Connection activity		
BER	16	0x00000010	Encoding and decoding of data, including ASCII and EBCDIC translations, if applicable		
FILTer	32	0x00000020	Search filters		
MESSage	64	0x00000040	Messaging subsystem activities and events		
ACL	128	0x00000080	Access Control List activities		
STATs	256	0x00000100	Operational statistics		
THREad	512	0x00000200	Threading activities		
REPLication	1024	0x00000400	Replication activities		
PARSe	2048	0x00000800	Parsing activities		
PERFormance	4096	0x00001000	Relational backend performance statistics		
Reserved	8192	0x00002000	Reserved		
REFErral	16384	0x00004000	Referral activities		
ERROr	32768	0x00008000	Error conditions		
SYSPlex	65536	0x00010000	Sysplex/WLM activities		
MULTIServer	131072	0x00020000	Multi-server activities		
LDAPBE	262144	0x00040000	Connection between a frontend and a backend		
STRBuf	524288	0x00080000	UTF-8 support activities		
TDBM	1048576	0x00100000	Relational backend activities (TDBM)		
SCHEma	2097152	0x00200000	Schema support activities (TDBM)		
BECApabilities	4194304	0x00400000	Backend capabilities		
CACHe	8388608	0x00800000	Cache activities		
ANY	2147483647	0x7FFFFFF	All levels of debug		
ALL	2147483647	0x7FFFFFF	All levels of debug		
Note: The minimum a	Note: The minimum abbreviation for each keyword is shown in uppercase letters.				

The debug level for the server can be set at a number of different times.

- The initial debug level is OFF.
- Prior to starting the server, the LDAP_DEBUG environment variable may be set. The server uses this value first. For example,

export LDAP DEBUG='ERROR+TRACE'

- When starting the server, the -d parameter may be specified. The debug level specified on this parameter either replaces, adds to or deletes from the preceding debug level. For example,
 - s ldapsrv,parms='-d ERROR'

replaces the current debug level that is either off or has been set by the LDAP_DEBUG environment variable with the new debug level of only ERROR.

- s ldapsrv,parms='-d +ERROR'

adds the ERROR debug level to the current debug level that is either off or has been set by the LDAP DEBUG environment variable.

- s ldapsrv,parms='-d -ERROR'

removes the ERROR debug level from the current debug level that is either off or has been set by the LDAP DEBUG environment variable.

 It is possible to change the debug level while the server is running whether it was started as a started task or from the shell. See "Dynamic debugging" on page 98 for more information.

All informational and error messages are printed to the started task output for the LDAP server.

When the LDAP server has been started and is ready, the message GLD0122I SLAPD is ready for requests.

is displayed.

"Running the LDAP server using data sets" discusses using a data set for the configuration file. In order to specify the configuration file as either a data set name or a DD name in SDSF, some special syntax is necessary.

In order to specify a full data set name, it may be necessary to be in the expanded input screen for SDSF. This is accomplished by entering a slash (/) in SDSF. On the expanded screen, it is then possible to specify a data set name for the configuration file. Assuming that the configuration file has been established in data set MYUSERID. SLAPD. CONF, the start command for the LDAP server in expanded input SDSF or the console would be:

```
s ldapsrv,parms='-f //'''MYUSERID.SLAPD.CONF''''
or, if additional parameters are desired:
```

```
s ldapsrv,parms='-f //'''MYUSERID.SLAPD.CONF''' -1 ldap://:999'
```

If a DD name, SLAPCONF, was established in the LDAPSRV PROC, as follows:

SLAPCONF DD DSN=MYUSERID.SLAPD.CONF,DISP=SHR

the LDAP server could be started from expanded input SDSF or the console by entering:

```
s ldapsrv,parms='-f //DD:SLAPCONF'
```

To stop the LDAP server in SDSF, enter:

/p ldapsrv

To stop the LDAP server from the operator's console, enter:

p 1dapsrv

This command causes the LDAP server to shut down.

Started task messages

The LDAP server writes messages to stdout and stderr. Messages sent to stdout and stderr appear in DD:SLAPDOUT in the provided JCL when running the LDAP server as a started task. SLAPDOUT appears in the started task log for the LDAP server and can be viewed through SDSF. See z/OS: SDSF Operation and Customization for information on how to use SDSF.

Running the LDAP server using data sets

Note: Using the LDAP configuration utility (Idapcnf) to configure your server creates all the necessary files in a partitioned data set.

The LDAP server, when run as a started task, accepts several of its files as data sets. Data set versions of the configuration files and envvars file are not shipped with the LDAP server, but can be created using the OGET command to copy the HFS versions of the files into data sets. (See z/OS: DCE Command *Reference* for information on the use of the **OGET** command.)

The default data set characteristics for record format and record length (V 255) which **OGET** will use when creating a new data set are not acceptable for JCL when submitting for batch processing. In order to avoid this, allocate the MYUSER.DSNTIJCL sequential data set to be fixed block 80 prior to performing the **OGET** operation.

A data set version of the DSNAOINI file needed for the TDBM backend can be created by copying and editing the default file provided by DB2. See step 4 on page 14. The DSNAOINI file can be specified either in the configuration file or in a DSNAOINI DD statement, or a DSNAOINI environment variable can be used. The **DD** statement takes precedence.

Note: Be sure that use of sequence numbers is turned off when editing this dataset.

Once the data set versions of these files are available, they can be specified in the LDAPSRV procedure. The configuration file can be specified using the CONFIG DD statement, the envvars file can be specified using the ENVVAR DD statement, and the DSNAOINI file can be specified using the DSNAOINI DD statement.

Setting up and running the LDAP server in the z/OS shell

A sample shell script is shipped for running the LDAP server in the shell. The sample shell script is located in /usr/lpp/ldap/sbin/slapd. The shell script needs to be modified to fit your environment. The PATH and STEPLIB variables should be set in the shell script. Ensure that /usr/sbin is added to the PATH environment variable. Also, be sure **STEPLIB** is set to *GLDHLQ*.SGLDLNK if the PDS has not been placed in LINKLIST, where GLDHLQ is replaced by the high level qualifier for the LDAP dataset. If your LDAP server is configured to encrypt passwords, the LIBPATH variable may need to be set. See "Installing OCSF and ICSF for password encryption" on page 16 for more information. The LDAP server writes messages and debugs to stdout and stderr. It is recommended that stdout and stderr be redirected and saved to capture any messages. Then, start the LDAP server by issuing: slapd >/tmp/slapd.log 2>&1 &

The same parameters described in "Setting up and running the LDAP server as a started task" on page 93 can be provided to the LDAP server when starting it from the z/OS shell.

It is also possible to define a separate user ID that will be used to run the LDAP server. See "Requirements for a user ID that runs the LDAP server" on page 35 for instructions on defining a separate user ID to run the LDAP server.

When started, the LDAP server will read an environment variable file. The default file is /etc/ldap/slapd.envvars. This default can be changed by setting the environment variable LDAP SLAPD ENVVARS FILE to the full path name of the desired environment variable file.

To stop the LDAP server in the z/OS shell, it is necessary to know its process ID. On any user ID that has a UID of 0, enter:

ps -ef | grep slapd

This will provide the process ID for the LDAP server. Next, enter:

kill -15 process-ID

where process-ID is the process ID of the LDAP server from the ps -ef command. This command will cause the LDAP server to shut down.

If multiple LDAP servers are running on the system, be sure to pick the correct server's process ID before entering the kill command.

Dynamic debugging

When the LDAP server is running as a started task or from the z/OS shell, it is possible to dynamically turn the debugging facility on and off. You can also replace the current debug levels, add to the current debug levels, or remove from the current debug levels. The following command can be sent to the LDAP server from the SDSF or the operator's console. Note that if the command is entered from SDSF, it must be preceded by a slash (/). In the command:

f ldapsrv,appl=debug=debug level

the debug level is a mask that specifies the desired debug level. (See Table 16 on page 95 for more information.) To send the same command to the LDAP server in the z/OS shell, it is necessary to know the job name assigned to the process by performing

d a,1

from SDSF or the operator's console and determining the name, which includes the user ID under which the LDAP server is running and a suffix. Once this name is found, use it to replace 1dapsrv in the command above. See Table 16 on page 95 for an explanation of the debug level values.

Debug information will be added to the output associated with the LDAP server.

To turn the debug tracing off, enter the same command providing the value zero (0) for debug_level.

Gathering trace records into the in-storage trace table

In addition to debugging, the LDAP server also keeps an internal trace table. A set of trace records is stored automatically (by default) in the in-storage trace table. Other traces, in addition to the default trace records, are controlled by a debug file. Use the _GLD_DEBUGFILE_NAME environment variable to specify the name of the debug file. The debug file contains entries which indicate what optional traces to include in the trace table. The size of the in-storage trace table can be set using the

GLD TRACE TABLE SIZE environment variable. The default size is 512K. You can specify the value in K, M, or bytes (for example, 1024K, 1M, 1048577). Note that once the trace table is full, the records begin to wrap (that is, the old records are overwritten by the new records).

You have the following options when gathering trace records for the LDAP server:

- · Determining which trace records go into the in-storage trace table
- Printing the trace table
- · Resetting the trace table

The command:

f ldapsrv,appl=debug

reads the debug file for the internal trace and resets all trace settings based on what is in the debug file. (Note that if the command is entered from SDSF, it must be preceded by a slash (/).) To send the same command to the LDAP server in the z/OS shell, it is necessary to know the job name assigned to the process. Do this by performing

d a,1

from SDSF or the operator's console and determining the name, which includes the user ID under which the LDAP server is running and a suffix. Once this name is found, use it to replace 1dapsrv in the command above.

Customizing the trace table

You can customize the trace entries produced in the in-storage trace table using the glddebug statements in the debug file. The **glddebug** statement has the following format:

```
glddebug=(c=type,p=prt,l=tLvl)cmt
```

where:

Is the trace type. Choose one of the following for *type*: type

- STORage
- LOCK
- ASYNCio
- LDAPDebug
- ALL

The uppercase shows the minimum abbreviation that can be used.

prt

Is the output direction. Choose one of the following for prt:

- m same place as in-storage trace
- y print to stderr
- **n** nothing in output

Is the trace level from 0-127. Each trace is identified by a type and level. The higher the number, tLvl the more traces that can be selected. For example, if you specify 120, 0-120 are selected.

Is a comment in the form: cmt

```
/* your comment */
```

Following is an example of using **glddebug**:

```
glddebug=(c=stor,p=m,l=127)
```

This statement would result in sending the output of all storage traces to the in-storage trace table.

Printing the trace table

Use the _GLD_TRACEFILE_NAME environment variable to specify the name of the file to which the trace table is printed. The file name defaults to /etc/ldap/qldtrace.output. If the trace file will not open, the trace file goes to stderr. Also, if there is an abend, the trace table is printed automatically.

To print the trace table, enter the following modify command:

f ldapsrv,appl=trace,print

Resetting the trace table

To reset the trace table (clear the trace table of its contents), enter the following modify command:

```
f ldapsrv,appl=trace,reset
```

Interaction between debug and in-storage trace

The debugging facility and the in-storage trace table interact as follows:

· The command

```
f ldapsrv,appl=debug=debug level
```

controls debug tracing. The command

f ldapsrv,appl=debug

(without the "=debug_level"), reads the debug file to start the internal trace.

A type setting of LDAPDebug on the glddebug statement directs debugs to the internal trace table.

Capturing performance information

The LDAP server provides a **query** command that is used to capture various performance information. The output of the query command goes to the system log and the job output. The syntax of the query command for LDAP is:

```
f ldapsrv,appl=query,report
```

Where report is the name of an LDAP performance report. The following is a list of all the available LDAP performance reports:

AIO This report shows the TCP/IP asynchronous socket calls made by the LDAP server and

the amount of queuing for incoming requests.

LOCKING This report shows the lock facility statistics. It includes the number of lock waits, the

average lock wait time, and the time threads sleep waiting for certain specific events.

THREADS This report lists the state of all the LDAP server threads.

ALL This option shows all of the reports described above.

Following is the sample output if you request an **AIO** report:

```
Asynchronous I/O Statistics
Service Thread Stack Size=76K Number of Pools=3
Pool: 0 Threads: 4 Pool: 1 Threads:
                                                                    10
        2 Threads:
                             0
Pool:
       SRB Accepts: 3 SRB Requests:
SRB Queued: 0 Pct. Queued:
Schedules: 8 Accepts:
Reads: 0 Readvs:
Recvs: 7 Recvfroms:
Writes: 0 Writevs:
Sends: 0 Cancelsockets:
                                                  Pct. Queued: 0.0%
                                                      Accepts :
                                                        Readvs :
                                                   Recvfroms:
                                                                               0
                                                      Writevs :
                                                                                0
                                                                                0
```

Activity Logging

The LDAP server can record server activity for the purpose of load analysis. The log file can contain information about operations handled by the server, messages generated by the server, and summary statistics. In general, each record consists of a time stamp followed by the record data. Operations are logged when they successfully begin processing. Optionally, a log record can be created when the operation completes processing. The following is an example of a typical operation log record followed by the ending log record.

```
Wed May 22 11:53:52 2002 Search: connid = 4, base = cn=monitor, filter = (objectclass=*)
Wed May 22 11:53:52 2002 End Search: connid = 4, base = cn=monitor, filter = (objectclass=*), rc = 0
```

Log records created for message logging appear in the log file with the timestamp followed by the message. The following is an example of a message log record.

```
Wed May 22 11:30:49 2002 GLD0210I Modify command has been processed successfully: LOG, MSGS
```

Summary records are created on an hourly basis as long as log records are being collected or when a modify command is processed that affects the log function. The summary log records contain information about the operations that the server has processed. The following is an example of the summary log records.

```
Wed May 22 11:45:25 2002 total operations started = 6
Wed May 22 11:45:25 2002 total operations completed = 6
Wed May 22 11:45:25 2002 total search entries sent = 2
Wed May 22 11:45:25 2002 total bytes sent = 897
```

```
Wed May 22 11:45:25 2002 total connections processed = 2

Wed May 22 11:45:25 2002 current connections = 2

Wed May 22 11:45:25 2002 connection high water mark = 2
```

- The location of the log file is controlled by the **logfile** configuration file option. The activity log can be written to either an HFS file or a MVS dataset. If the log file is an MVS dataset, it must be created (allocated) prior to its use by the LDAP server. It is recommended that when an HFS file is desired that the file specification be fully qualified. The following is an example of a logfile option specifying an HFS file.
- | logfile /etc/ldap/gldlog.output
- /etc/ldap/gldlog.output is also the default location of the log file
- The MVS dataset can be specified through either a ddname or as a specific dataset. The following is an example for both methods.
- | logfile dd:logout
- | logfile //mysys.ldap.actlog
- The default data collection setting is to collect no data. The default is modified by either environment variables or through the specification of operator modify commands. The environment variables are read as the server starts up while the modify commands can be specified once the server has started.
- A modify command can be sent to the LDAP server from the SDSF or the operator's console. Note that if the command is entered form SDSF, it must be preceded by a slash (/). In the command,
- f ldapsrv,appl=log,setting
- the *setting* is the modification to the log collection. To send the same command to the LDAP server in the z/OS shell, it is necessary to know the job name assigned to the process. Do this by performing
- from SDSF or the operator's console and determine the name which includes the user ID under which the LDAP server is running and a suffix. Once this name is found, use it to replace Idapserv in the command above.
 - The **GLDLOG_OPS** environment variable or the modify command controls which operations generate log entries. Prior to specifying an operation setting, no operation logging is performed. The settings are:
- writeops indicates that log entries are to be created at the start of add, delete, modify, modrdn and extended operations.
- **allops** indicates that log entries are to be created as for writeops and in addition, search and compare operations are to be logged.
- **summary** indicates that only hourly summary statistics are to be logged. As long as any logging is being collect, summary data is produced on an hourly basis.
- The **GLDLOG_TIME** environment variable or the modify command controls whether log entries are generated when logged operation ends. The settings are:
- time indicates that ending logs are to be created for all operations that are being logged.
- notime indicates that ending logs are not to be created for operations that are being logged. This is the
 default prior to specifying the GLDLOG_TIME environment variable or by issuing the modify command
 to set the operation end time control.
- The **GLDLOG_MSG** environment variable or the modify command controls whether log entries are generated when messages are created by the LDAP server. The settings are:
- msgs indicates that messages generated by the LDAP server are to be written to the activity log in addition to the normal target.

 nomsgs indicates that messages generated by the LDAP server are not to be written to the activity log. This is the default prior to specifying the GLDLOG MSG environment variable or by issuing the modify command to set the message control.

As the log entries are produced some buffering of the output is performed by the system. The buffers are flushed before the server shuts down. However, you can force the server to flush the buffers by entering a modify command:

```
f ldapsrv,appl=log,flush
```

The server can be told to stop collecting activity data by entering a modify command:

f ldapsrv,appl=log,stop

Using the LDAP server for PC callable support

When the LDAP server comes up, it tries to enable any interfaces that are configured to it. If at least one interface comes up, the LDAP server remains up. If you are using an LDAP server for program call (PC) callable support, it is recommended that you use it only for that type of interface and no other interfaces. At most one LDAP server in a system can activate PC callable support. If an LDAP server tries to initialize PC callable support after another LDAP server has attempted (successfully or unsuccessfully) to initialize PC callable support, the initialization fails. The previous server that has locked access to the PC callable support must be shut down before another server can attempt to run with the PC callable interface.

Verifying the LDAP server

The following examples show how you can verify your LDAP server using the Idapsearch tool. Note that you can use any LDAP client to do this.

Verifying TDBM

In the command below, substitute the suffix value from your configuration file for the -b parameter. The command can be run multiple times to verify each suffix defined in the configuration file.

```
ldapsearch -h 127.0.0.1 -b "o=Your Company" "objectclass=*"
```

Verifying SDBM

For SDBM, you must bind with a valid RACF-style DN to perform the search. Substitute a RACF ID of your choice in the racfid portion of the DN on the -D and the -b parameters below. Also, substitute your SDBM suffix in the DN on the -D and -b parameters. The RACF password for the user ID used in the **-D** parameter must be specified in the **-w** parameter.

```
ldapsearch -h 127.0.0.1 -D racfid=IBMUSER,profiletype=user,cn=myRacf
  -w password_for_IBMUSER -b racfid=IBMUSER,profiletype=user,cn=myRacf "objectclass=*"
```

Both of the previous Idapsearch examples assume a default port of 389. If your port is not 389, use the **-p** parameter to specify the correct port.

Be sure to substitute the correct TCP/IP host name or TCP/IP address for the 127.0.0.1 after the -h parameter. The **-b** parameter specifies the starting point for the search. The use of the quotation marks around the filter prevents the asterisk (*) from being interpreted by the shell.

Note that this can be done from TSO as well, substituting **LDAPSRCH** for **Idapsearch**.

See *z/OS*: Security Server LDAP Client Programming for more information about **Idapsearch**.

Chapter 10. Migrating to a z/OS LDAP server

This chapter discusses migration issues. Your plan for migrating to a z/OS LDAP server should include information from a variety of sources. These sources of information describe topics such as coexistence service, hardware and software requirements, migration actions, and interface changes.

The following documentation, which is supplied with your product order, provides information about installing your z/OS system. In addition to specific information about the LDAP server, this documentation contains information about all of the z/OS elements.

• z/OS: Planning for Installation

This book describes the installation requirements for z/OS at a system and element level. It includes hardware, software, and service requirements for both the driving and target systems. It also describes any coexistence considerations and actions.

• z/OS: Program Directory

This document, which is provided with your z/OS product order, leads you through the specific installation steps for the LDAP server and the other z/OS elements.

• ServerPac: Installing Your Order

This is the order-customized, installation book for using the ServerPac Installation method. Be sure to review "Appendix A. Product Information," which describes data sets supplied, jobs or procedures that have been completed for you, and product status. IBM may have run jobs or made updates to PARMLIB or other system control data sets. These updates could affect your migration.

Within this chapter, you can find information about the specific updates and considerations that apply to this release of z/OS LDAP.

· "Steps for migrating"

This section identifies tasks that must be done for any migration. These tasks are not associated with a specific new function, but must be completed to ensure successful migration to the new release.

"Actions required for migrations" on page 104

This section describes the common activities and considerations to consider depending on the particular release you are migrating from.

"Migration roadmap" on page 110

This section identifies the migration paths that are supported with the current level of the LDAP server. It also describes the additional publications that can assist you with your migration to the current level.

"z/OS V1R4 overview" on page 111

This section describes the specific updates that were made to LDAP for the current release. For each item, this section provides an overview of the change, a description of any migration and coexistence tasks that may be considered, and where you can find more detailed information in the z/OS LDAP library or other element libraries.

"Summary of interface changes" on page 123

This section provides a summary of the changes that are made to the LDAP server user and programming interfaces.

Steps for migrating

The recommended steps for migrating to a new release of z/OS LDAP are:

- Become familiar with the supporting migration and installation information for the release.
 Determine what updates you need for: products that are supplied by IBM, system libraries, and non-IBM products. Review z/OS: Planning for Installation and z/OS: Introduction and Release Guide for information about the LDAP server and other z/OS elements.
- 2. Develop a migration plan for your installation.

When planning to migrate to a new release of the LDAP server, you must consider high-level support requirements, such as machine and programming restrictions, migration paths, and program compatibility.

3. Obtain and install any required program temporary fixes (PTFs) or updated versions of the operating system.

Call the IBM Software Support Center to obtain the preventive service planning (PSP) upgrade for the z/OS LDAP server, which provides the most current information about PTFs for the z/OS LDAP server. Check RETAIN® again just before testing the LDAP server. For information about how to request the PSP upgrade, refer to the z/OS: Program Directory. Although the z/OS: Program Directory contains a list of the required PTFs, the most current information is available from the IBM Software Support Center.

- 4. Install the product using the z/OS: Program Directory or the ServerPac: Installing Your Order documentation.
- 5. Contact programmers who are responsible for updating applications at your installation. Verify that your installation's applications will continue to run, and, if necessary make changes to ensure compatibility with the new release.
- 6. Use the new release before initializing major new function.
- 7. If necessary, customize the new function for your installation.
- 8. Exercise the new functions.

Actions required for migrations

The following sections describe common activities and considerations that are typically required (or should be considered) whenever you migrate from a previous release of the LDAP server to the current release of the LDAP server. In this chapter, all references to OS/390 V2R10 also apply to z/OS V1R1 and all references to z/OS V1R2 also apply to z/OS V1R3.

Note: With APAR OW50971 for OS/390 V2R10 and z/OS V1R1, the complete LDAP product at the z/OS V1R2 level was rolled back to OS/390 V2R10 and z/OS V1R1. If this APAR is applied to an OS/390 V2R10 or z/OS V1R1 system, all functionality of the z/OS V1R2 LDAP Server becomes available on OS/390 V2R10 or z/OS V1R1. If this rollback APAR has been applied to your OS/390 V2R10 or z/OS V1R1 system, you should follow sections marked with z/OS V1R2.

After each heading in this section, a table indicates the OS/390 or z/OS LDAP releases you could be migrating from, showing an **X** for each release the particular section pertains to. If, for example, you are migrating from OS/390 V2R10, the sections that show an X under OS/390 V2R10 and z/OS V1R1 apply to your installation and should be considered.

Obtaining the DB_VERSION setting for RDBM

	OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
ļ	X	X

Some of the steps in the following migration actions depend on the DB_VERSION setting for existing
RDBM DB2 tables. The DB_VERSION can be obtained using the following SQL operation entered through
SPUFI:

SELECT DB VERSION FROM USERID.LDAP NEXT EID

Where *USERID* is the value of the **dbuserid** option in the configuration file.

Obtaining the DB_VERSION setting for TDBM

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	X

Some of the steps in the following migration actions depend on the **DB VERSION** setting for existing TDBM DB2 tables. The DB VERSION can be obtained using the following SQL operation entered through SPUFI:

| SELECT DB VERSION FROM USERID.DIR MISC

Where *USERID* is the value of the **dbuserid** option in the configuration file.

Coexistence and migration with previous releases (RDBM)

RDBM has been removed. All customers using RDBM must migrate to TDBM.

Because TDBM has a different DB2 table design than RDBM, existing RDBM data must be unloaded and reloaded when migrating to TDBM. For the same reason, DB2 data sharing between TDBM and existing RDBM tables is not possible.

Migrating existing RDBM data to a TDBM database

- 1. Create a backup of the existing RDBM directory data using native DB2 utilities.
- 2. The RDBM database must be unloaded on the existing release prior to moving to z/OS V1R4. Unload the RDBM directory information using **db2ldif**. As an example:

db2ldif -f /etc/ldap/slapd.conf -o /tmp/directorydata.unload

assuming this command was run from the z/OS shell prompt.

- 3. Obtain the **DB VERSION** setting for the existing RDBM DB2 tables (see "Obtaining the DB VERSION setting for RDBM" on page 104 for more information). If the **DB_VERSION** is 8.0, skip to step 8 on page 106. Otherwise, evaluate all comments (lines beginning with a pound sign (#)) that begin with AF, AP, DF, or DP in the LDIF file created in the previous step. If the attribute value or distinguished name described in this comment is acceptable to you, no further updates are necessary for that attribute value. If the attribute value is not acceptable to you, modify the uncommented LDIF line below the comments to contain an attribute value acceptable to you.
- 4. Create the TDBM database and table spaces in DB2. See "Creating the DB2 database and table spaces for TDBM" on page 40 for these instructions.
- 5. Make appropriate configuration changes. See "Configuration file checklist" on page 52 for information on the TDBM configuration options. The RDBM section and the includes for schema files should be removed.
- 6. The schema definitions associated with the TDBM backend conform to the LDAP Version 3 (V3) protocol. The schema files that are shipped with z/OS LDAP to be used with TDBM are a superset of the schema definitions shipped for use with RDBM. The schema files shipped for TDBM are in the LDAP V3 protocol format while the schema files shipped for RDBM were in LDAP V2 protocol format. If you have modified the LDAP V2 protocol schema files used with RDBM, corresponding changes may need to be made to the LDAP V3 protocol format schema files used with TDBM. If you have added schema definitions, then corresponding schema definitions may need to be added to the LDAP V3 protocol schema files. See "Migrating the schema from RDBM to TDBM" on page 176 for more
 - 7. Load the data from the LDIF file just created using either the **Idif2tdbm** utility or the **Idapadd** utility. See "Idif2tdbm program" on page 127 for details, including instructions on determining which utility is best suited for your installation.

- 8. Optionally, remove the old DB2 tables. To do this, remove the tablespaces for the RDBM database. The following command should be repeated for each tablespace associated with the RDBM database. Examine the SPUFI used to create the RDBM database to determine the tablespace names.
 - DROP TABLESPACE <tbspname>

This step is optional because RDBM and TDBM databases can coexist.

Migrating RDBM to TDBM: entryOwner and aclEntry attribute value changes

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	X

With RDBM, it is possible to add entryOwner and aclEntry attribute types where the distinguished name (DN) portion of the attribute value contains an incorrectly formatted DN. For example, unescaped special characters (such as the # sign) are allowed. TDBM does not support incorrectly formatted DNs in aclEntry and entryOwner attribute values. Attribute values of this type containing correctly formatted DNs can be migrated to TDBM without change. Attribute values of this type containing incorrectly formatted DNs must be manually corrected prior to being loaded into TDBM. To date, this problem has only been encountered with SDBM DNs present in RDBM aclentry and entryOwner values that were being migrated to TDBM. See IETF RFC 2253, Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names for rules that define a correctly formatted DN.

Migrating RDBM to TDBM: timestamp changes

	OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
1	X	X

If you are migrating from an RDBM backend to a TDBM backend, the timestamps are different. Timestamp attribute values generated by the RDBM backend (modifytimestamp and createtimestamp) were of the following format:

modifytimestamp=2000-05-05 10:30:56.128806

This value represents the "local time". In this example, the local time is with respect to eastern standard time timezone. (Note that "daylight savings time" is in effect.)

Timestamp attribute values generated by TDBM are in the Generalized Time format with respect to GMT time. In TDBM, the previously listed timestamp would be:

modifytimestamp=20000505143056.128806Z

Migrating from port, securePort, and security to listen option

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	

The port, securePort, and security configuration options have been deprecated by the listen option in z/OS V1R2. In order to migrate from an earlier configuration file to a current one which uses the **listen** option, check your current security option setting. If security is set to:

none or nossl

Add the following to the configuration file:

listen ldap://:port

The *port* is equal to the **port** option setting (if specified).

sslonly

Add the following to the configuration file:

listen ldaps://:port

The *port* is equal to the **securePort** option setting (if specified).

ssl Add the following to the configuration file:

listen ldap://:port1
listen ldaps://:port2

The *port1* is equal to the **port** option setting (if specified) and *port2* is equal to the **securePort** option setting (if specified).

The **security**, **port**, and **securePort** options can now be removed from your configuration file.

Note: If **port** or **securePort** are not specified in your current configuration file, then the default of 389 is used for the nonsecure port and 636 is used for the secure port.

Migrating from slapd command-line arguments -p and -s to -l

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	

The **-p** port and **-s** secureport options on the LDAP server command line are no longer evaluated by the LDAP server when one or more **listen** options are specified either on the command line or in the configuration file. It is strongly recommended that you migrate to the new **-l** listen option that is introduced in z/OS V1R2. Following is an example of how to migrate from the **-p** and **-s** options specified on the command line to the one or more equivalent **-l** options. Given the following LDAP server command:

slapd -p port1 -s port2

- Migrate the port (-p) command line option to a listen (-l) command line option:
 - -1 ldap://:port1
- Migrate the secure port (-s) command line option to a listen (-l) command line option:
 - -1 1daps://:port2
- Remove the -p and -s options from the command line and replace with the -I options.
- The migrated example is now:

slapd -l ldap://:port1 -l ldaps://:port2

Migrating from maxThreads and waitingThreads to commThreads option

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	

The **maxThreads** and **waitingThreads** options are no longer evaluated by the LDAP server. Both of these options have been replaced by the **commThreads** option that is now available. These options should be removed from the configuration file prior to running the LDAP server. If they are not removed, messages will be issued and the options will be ignored. The **commThreads** option is used to specify the number of

threads created at LDAP server initialization. Previously, maxConnections and maxThreads had to be the same. This is no longer a requirement with maxConnections and commThreads.

It is recommended that commThreads be set to approximately two times the number of CPUs that are running in your LPAR. However, this is a general rule depending upon the IO activity that your LDAP server experiences.

Migrating to the updated maxConnections option

	OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
ļ	X	

In z/OS V1R2, the maxConnections option should be set to the expected number of concurrently connected clients that the LDAP server will allow. Previously, this option was limited by maxThreads, which has been deprecated in z/OS V1R2. Prior to z/OS V1R2, both maxConnections and maxThreads were strongly suggested to be the same value. The default value for maxConnections has been updated from 75000 to the system maximum. The minimum value for maxConnections has been changed from 10 to 30. It is suggested that to migrate from a previous version of the configuration file, remove the maxThreads option and add ten to the current value of maxConnections. Refer to page 62 for further explanation on maxConnections.

APF authorizations

	OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
ļ	X	

Prior to z/OS V1R2 of the LDAP server, it was not always necessary to APF-authorize the LDAP server, depending upon which backends were configured. Beginning in z/OS V1R2, the LDAP server PDS must be APF-authorized no matter which backends are configured. See "Protecting the environment for the LDAP server" on page 37 for more information.

Schema migration

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	X

Schema updates have been shipped as part of service for the OS/390 and z/OS LDAP servers. The application of these changes as found in the /usr/lpp/ldap/etc/schema-updates.ldif file by the LDAP administrator is recommended prior to starting the z/OS V1R4 LDAP server. See the schema-updates.ldif file and "Updates to the schema" on page 160 for additional information. All future schema service will assume that these updates have been applied to customer schemas.

Migrating TDBM databases

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and z/OS V1R3
X	X

In order to use the new access control grant, deny, and attribute-level permissions capabilities provided in z/OS V1R4, the TDBM database version (as defined in the DB2 tables allocated for the TDBM database backend) must be set appropriately. For database instances created using z/OS V1R4, the database

version will be set to 2.0 as part of first starting up either the LDAP server or Idif2tdbm programs. For database instances created using an earlier release, the TDBM database version must be updated in order to use the new access control features.

Releases of the z/OS Security Server LDAP server prior to z/OS V1R4 do not contain support for the enhanced access control capabilities. Thus, if you are running the LDAP server in a sysplex environment with mixed levels of systems, you cannot use the enhanced access control support until all LDAP servers accessing the same TDBM database instance are running at the z/OS V1R4 level. Once all systems on which LDAP servers are running that are sharing the same DB2 tables are running z/OS V1R4, the database version can be set to 2.0 and the enhanced access control capabilities can be used.

In order to change the database version of an existing TDBM database instance that was created using an LDAP server from a release prior to z/OS V1R4, run the following SQL using either the SPUFI facility or suitable program (for example, DSNTIAD):

UPDATE USERID.DIR MISC SET DB VERSION='2.0';

where USERID is set to be the value of the dbuserid LDAP server configuration file setting for the TDBM backend instance.

To verify that the enhanced access control capabilities are enabled, examine the LDAP server or **Idif2tdbm** program output log for message:

| GLD3135I Grant/Deny ACL support is enabled below suffixes: suffix-list.

Fallback procedures in case a prior release of the LDAP server must be run

If it becomes necessary to fall back to a prior release of the LDAP server, it is possible to set the TDBM database version back to '1.0'. By doing this, enhanced access control capabilities will be turned off and any enhanced access controls that were set up while the TDBM database version was set to '2.0' will be ignored for access control checking. Also note that any entries which contain enhanced access controls will not be modifiable while running in this fall back mode. To set the TDBM database version back to '1.0', run the following SQL:

UPDATE USERID.DIR MISC SET DB VERSION='1.0';

To verify that the enhanced access control capabilities are disabled, examine the LDAP server or Idif2tdbm program output log for message:

GLD3136I Grant/Deny ACL support is not enabled below suffixes: suffix-list.

If the entry containing the enhanced aclEntry information must be updated, the z/OS V1R4 LDAP server must be used. To make this update, stop the LDAP server, modify the DB VERSION value back to 2.0, and start the z/OS V1R4 LDAP server. Verify that message GLD3135I appears in the server log. Modify the entry containing the enhanced aclEntry information. Once the updates are complete, stop the z/OS V1R4 LDAP server. To run a prior release of the LDAP server, once again follow the fallback procedures to modify the DB VERSION value and start the LDAP server

Note: Since enhanced access controls allow for deny access control specifications, there exists the possibility that users who could not access information when the enhanced access controls are supported will be able to access information when the enhanced access controls are ignored (while running in fallback mode). If enhanced access controls are being used and fallback is necessary, it is recommended that this possibility be examined prior to running in a fallback mode.

Removing schema file include statements

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and V1R3
X	X

Remove the include statements in your configuration file for any files that contain attribute or objectclass lines. These definitions were needed in OS/390 V2R10 and z/OS V1R2 for SDBM and RDBM. The attribute and objectclass configuration options are no longer supported. If you have defined your own attributes and objectclasses for use with RDBM, refer to "Schema migration" on page 108 for information about converting your definitions for use with TDBM.

Removing verifySchema configuration option

OS/390 V2R10 and z/OS V1R1	z/OS V1R2 and V1R3
X	X

Remove the verifySchema configuration option statement from your configuration file. It is no longer supported.

Migration roadmap

This section describes the migration paths that are supported by the current release of the LDAP server. It also provides information about how you can obtain the LDAP server migration information from previous releases.

z/OS V1R4 update summary

The following table summarizes the updates introduced to the z/OS LDAP server in z/OS Version 1 Release 4 (z/OS V1R4). If you are migrating from z/OS V1R3, z/OS V1R2, z/OS V1R1, or OS/390 V2R10, you should review the information in the detailed section for each item.

Table 17. Summary of LDAP server updates for z/OS V1 R4

1	For Information About:	Refer to Page:
	Modify DN	183
	CRAM-MD5 and DIGEST-MD5 authentication	243
	Transport Layer Security (TLS) support	43
	ACL enhancements	249
	SDBM enhancements	39
	LDAP IBM-entryuuid support	141
	Activity logging	93
	Abandon support	297
	RDBM removal	105
	Monitor support	297

z/OS V1R2 update summary

The following table summarizes the updates introduced to the z/OS LDAP server in z/OS Version 1 Release 2 (z/OS V1R2). If you are migrating from z/OS V1R3, z/OS V1R2, z/OS V1R1, or OS/390 V2R10,

you should review the information in the detailed section for each item.

Table 18. Summary of LDAP server updates for z/OS V1R2

For Information About:	Refer to Page:
Kerberos authentication	225
Native authentication	233
RACF connect profile support and other SDBM enhancements	205

Table 18. Summary of LDAP server updates for z/OS V1R2 (continued)

For Information About:	Refer to Page:
Support for LDAP extended operations that retrieve Policy Director data	247
Extended operations: GetDnForUserid and GetPrivileges	435
Additional configuration options: commThreads, idleConnectionTimeout, krbIdentityMap, krbKeytab, krbLDAPAdmin, listen, nativeAuthSubtree, nativeUpdateAllowed, pcThreads, serverKrbPrinc, supportKrb5, and useNativeAuth	53
maxConnections migration updates	62
Deprecated configuration options: waitingThreads, maxThreads, port, securePort, and security	73
Specifying a debug level as a mask	95
Scalability improvements	62

z/OS V1R1 update summary

No new function was introduced to the z/OS LDAP server for this release.

OS/390 V2R10 update summary

Table 19 summarizes the updates that were introduced to the OS/390 LDAP server in V2R10 (OS/390 V2R10).

Table 19. Summary of LDAP server updates for OS/390 V2R10

For information about:	Refer to page:
Schema publication and update	159
TDBM backend based on DB2, including new load and unload facilities (ldif2tdbm and tdbm2ldif)	"Creating the DB2 database and table spaces for TDBM" on page 40
	127 (Idif2tdbm) 137 (tdbm2ldif) 259
LDAP server configuration utility (Idapcnf)	17
Additional configuration options: attrOverflowSize and sslCertificate	53
New ACL administration methods when using TDBM	249

z/OS V1R4 overview

This section describes the new and changed LDAP server functions introduced for z/OS Version 1 Release 4 (V1R4). The information about each item includes:

- Description
- Summary of the LDAP server tasks or interfaces that may be affected
- Coexistence considerations, if any, that are associated with the item
 - Migration procedures, if any, that are associated with the item
- References to other publications that contain additional detailed information.

Modify DN

Description

Following is a description of each Modify DN enhancement:

- A Modify DN operation can now be accompanied by the new superior parameter. This parameter specifies the DN that will become the immediate superior of the renamed entry.
- A non-leaf node can now be the target of a Modify DN operation.
- · A subtree move can be accomplished by specifying a New Superior parameter when the target of the Modify DN operations is a non-leaf node.
- The new IBMModifyDNRealignDNAttributesControl can accompany a Modify DN operation. This control will cause the server to search for all attributes of DN syntax whose values match the old DN value. These values will then be updated to reflect the new DN that was created in the Modify DN operation.

What this change affects

Modify DN operations will now affect Replication. For more information see Chapter 17, "Modify DN Operations" on page 183.

Dependencies

None.

Coexistence considerations

Replication of enhanced Modify DN operations only works with z/OS V1R4 and above LDAP servers. For more information see Chapter 17, "Modify DN Operations" on page 183.

Migration tasks

None.

For more information

See Chapter 17, "Modify DN Operations" on page 183.

CRAM-MD5 and DIGEST-MD5 authentication

Description

The z/OS LDAP server has been enhanced to allow CRAM-MD5 and DIGEST-MD5 authentication.

What this change affects

Allows SASL bind mechanisms of CRAM-MD5 and DIGEST-MD5.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

None.

For more information

See Chapter 21, "CRAM-MD5 and DIGEST-MD5 Authentication" on page 243.

Transport Layer Security (TLS) support

Description

The z/OS LDAP server and client APIs provide support through System SSL for Transport Layer Security (TLS) protocol protection of client/server communication. Transport Layer Security is based upon Secure Sockets Layer (SSL) Security version 3. The differences between TLS v1.0 and SSL v3.0 protocol are not dramatic, but they are significant enough that TLS v1.0 and SSL v3.0 do not inter operate. The z/OS LDAP server and client continues to support SSL v3.0 protocol protection through System SSL. System SSL determines the protocol to use during the secure protocol handshake performed by the peers.

The z/OS LDAP server supports the Start TLS extended operation as described in RFC 2830. This extended operation allows a client communicating over a non-secure connection to change communication to secure communication protected by SSL/TLS. The client may later terminate secure communication while maintaining the connection to the server resulting in non-secure communication between the server and the client.

What this change affects

This change allows the z/OS LDAP server and client APIs to use Transport Layer Security for secure communication. In addition, the LDAP server supports the Start TLS extended operation to change a non-secure connection to a secure connection. The z/OS client APIs do not support the Start TLS extended operations.

Dependencies

z/OS V1R4 level of System SSL is required.

Coexistence considerations

None.

Migration tasks

None.

For more information

See "Setting up for SSL/TLS" on page 43.

ACL enhancements

Description

Additional capabilities have been added to the TDBM backend database to support grant/deny as well as attribute-level access control permissions. Access control lists (ACLs) can now contain permission clauses which explicitly deny access to information. In prior releases, access could only be granted, with access to all other information being denied implicitly. ACLs can now be set up to grant or deny access to individual attribute types. In prior releases, access could only be granted to attribute access classes (groups of attributes).

What this change affects

This change affects the format of values of the aclEntry attribute. The changed format is a super set of the format that was supported in prior releases. Access control lists created using prior releases continue to work as before.

Dependencies

None.

Coexistence considerations

In order to use this support, all LDAP servers that are using the same set of DB2 tables (TDBM backend database instance) must be capable of interpreting the additional format for values of the aclEntry attribute. When running in a sysplex or multi-server mode, all LDAP servers must be running at the z/OS V1R4 level before this support can be enabled. This support is enabled by modifying the DB VERSION value for the TDBM backend database instance. New TDBM backend database instances, created using the z/OS V1R4 LDAP server, will initialize the DB_VERSION value to allow the enhanced controls to be used in the newly defined TDBM backend databse instance. In order to use such a TDBM access backend database instance in a sysplex or multi-server mode with z/OS LDAP servers from previous releases, the DB VERSION value must be set to 1.0.

This support also requires the coexistence PTF for OS/390 V2R10 and z/OS V1R2 (OW54426).

Migration tasks

Refer to "Migrating TDBM databases" on page 108 for detailed information on modifying the DB_VERSION value for the TDBM backend database instance.

For more information

See Chapter 23, "Using access control" on page 249 for more information on using the new grant/deny and attribute-level access control capabilities.

SDBM enhancements

Description

SDBM has been enhanced to support:

- 1. Setting and displaying the values of the new EIM segment in the RACF user profile.
- 2. Specifying SHARED and AUTOUID or AUTOGID when setting the UID or GID value in the OMVS segment in the RACF user or group profile.
- 3. Searching for all the RACF users or groups with a specified OMVS UID or GID value.
- 4. Using an internal schema instead of defining the attributes and object classes in external schems files.
- SDBM has also changed the reason code that it returns when a client attempts to bind using a user who doesn't exist, a user who is not fully defined, or an incorrect password. All these cases now return reason code R000104. To cover all these cases, the text for reason code R000104 has been changed to:
- R000104 The password is not correct or the user id is not completely defined (missing password or uid).

What this change affects

The new RACF control of shared UID/GID values can now be used when adding or modifying a user or group using SDBM. The reason code returned when binding to SDBM has changed as described above.

Dependencies

Enhancements 2 and 3 require that RACF be at AIM Stage 2 and require additional RACF setup. See the *z/OS: RACF Security Administrator's Guide* for more information.

Coexistence considerations

In previous releases, binding to SDBM with a user who doesn't exist or who is not fully defined returned:

R000103 The user id is not completely defined - missing password or uid.

It now returns R000104 as described above.

Migration tasks

The use of an internal schema by SDBM instead of external schema files has several effects:

- 1. SDBM no longer uses external schema files. Any **include** options for schema files should be removed from the configuration file.
- An SDBM suffix should not contain an alias for an attribute. For example, an SDBM suffix cannot use the **surName** attribute (it can use the **sn** attribute instead). Also, an SDBM suffix can contain a case-sensitive attribute in the suffix, but SDBM ignores case when processing the suffix. Applications that check the reason code returned for when binding to SDBM may need to be modified to accommodate the reason code change described above.

For more information

See Chapter 18, "Accessing RACF information" on page 205.

LDAP IBM-entryuuid support

Description

The z/OS LDAP server provides support for **IBM-entryuuid**.

What this change affects

Adds an attribute which contains a unique identifier to every LDAP entry.

Dependencies

None.

	Coexistence considerations Only works with OS/390 V1R2 with the coexistence APAR OW54426 and above z/OS LDAP servers.
	Migration tasks None.
	For more information See Chapter 12, "Running and using the LDAP entry UUID utility" on page 141 and the serverEtherAddr option on page 66.
	Activity logging
	Description The z/OS LDAP server provides a means of collecting server activity information in a log file. The log file may be an HFS file or an MVS dataset.
	What this change affects A new option has been added to the configuration file to allow specification of the log file. New modify commands are supported by the z/OS LDAP server to control generation of the log file.
	Dependencies None.
	Coexistence considerations None.
	Migration tasks None.
	For more information For more information, see "Activity Logging" on page 100
	Abandon support
	Description The Abandon operation was supported in previous releases of the LDAP server. Hoever, all operations on a single connection were processed synchronously. Therefore, the abandon operation had no effect on the processing or previous operations by the LDAP server. In z/OS V1R4, the LDAP server reads additional operations as they arrive as long as the connection is not a secure connection and the previous operation is not bind, unbind, or extended operations.
	What this change affects None.
	Dependencies None.
	Coexistence considerations None.
	Migration tasks None.
	For more information None.

RDBM removal

Description

Support for the RDBM backing store is removed. Customers using this DB2-based backing store must migrate from use of RDBM to use of TDBM, which is also DB2 based.

What this change affects

TDBM allows storage of larger amounts of directory data and provides improved performance over RDBM.

Dependencies

None.

Coexistence considerations

Since RDBM is not available in z/OS V1R4 LDAP, customers wishing to add z/OS V1R4 to a sysplex where they are already using the LDAP server and sharing an RDBM database across the sysplex must

migrate their shared database to TDBM across the sysplex before adding z/OS V1R4 into the sysplex.

TDBM is available beginning in OS/390 V2R10 and is included in all z/OS releases. It is not possible to

share data across a sysplex between an RDBM backing store and a TDBM backing store.

Migration tasks

To migrate from RDBM to TDBM, customers must unload their RDBM backing store to LDIF and reload that LDIF into a newly-defined TDBM backing store.

For more information

See "Coexistence and migration with previous releases (RDBM)" on page 105.

Monitor support

Description

The LDAP server provides the capability to retrieve certain statistics about its operations through an LDAP server using a base of cn=monitor.

What this change affects

Provides additional monitoring capability.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

None.

For more information

See "Monitor Support" on page 297.

TDBM schema

Description

Additional attribute types and object classes are added to the TDBM schema by applying the updates from the schema-update.ldif file shipped with this release. Enhanced schema migration instructions are now available.

What this change affects

The TDBM schema is enhanced by the addition of the attribute types and object classes. Future updates to the schema will be simplified.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

Current users if the TDBM backend who have not yet applied APAR OW53447 need to perform the schema upgrade tasks defined in Chapter 16, "LDAP directory schema for TDBM" on page 159.

For more information

See "Schema migration" on page 108 for additional information about migrating the schema.

LDAP server message enhancements

Description

Changed the severity level to be consistent with other z/OS product message severities. Modified the LDAP server message numbers to correctly reflect the severity level documented in the message

description.

What this change affects

The message numbers have changed.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

Use of message numbers in automation of LDAP events should be reviewed and modified to be consistent with the new message number values..

For more information

See Chapter 28, "LDAP server messages" on page 317 for more information.

z/OS V1R3 overview

No new function was introduced to the z/OS server for this release.

z/OS V1R2 overview

This section describes the new and changed LDAP server functions introduced for z/OS V1R2 (V1R2). The information about each item includes:

- Description
- Summary of the LDAP server tasks or interfaces that may be affected
- · Coexistence considerations, if any, that are associated with the item
- · Migration procedures, if any, that are associated with the item
- · References to other publications that contain additional detailed information.

Kerberos authentication

Description

The z/OS LDAP server has been enhanced to allow Kerberos Authentication.

What this change affects

Allows a SASL bind mechanism of GSS API.

Dependencies

Security Server Network Authentication and Privacy Service must be installed and configured.

Coexistence considerations

None.

Migration tasks

None

For more information

See Chapter 19, "Kerberos authentication" on page 225.

Native authentication

Description

The z/OS LDAP server has been enhanced to allow authentication to the Security Server through the TDBM backend.

What this change affects

Allows authentication to be performed by the Security Server.

Dependencies

A Security Server must be installed and configured.

Coexistence considerations

None.

Migration tasks

None.

For more information

See Chapter 20, "Native authentication" on page 233.

SDBM enhancements

Description

Following is a description of each SDBM enhancement:

- RACF manages group membership and maintains a set of connection information about each member of a group. SDBM can now be used to administer this connection information.
- Support has been added for the LNOTES, NDS, KERB, and PROXY segments in the RACF user profile. Values for these segments can be set and displayed in an SDBM user profile entry. SDBM can also be used to administer RACF universal groups.
- New application identification search filters have been added to SDBM. These filters allow an application to retrieve the entire entry of the RACF user corresponding to a given LNOTES short name, NDS user name, KERB principal name, or OMVS UID. The entire entry of the RACF group corresponding to a given OMVS GID can also be retrieved.

What this change affects

With the LDAP connect profile entry you can add or remove users from a group and you can display or modify the connection information. Connection information is represented as another subtree under the suffix and can be returned in subtree searches beginning at the suffix.

The user profile entry can now include additional object classes and attributes representing the LNOTES. NDS, KERB, and PROXY segments. The group profile entry can now include an attribute indicating a universal group.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

None.

For more information

See Chapter 18, "Accessing RACF information" on page 205.

Extended operations for accessing Policy Director data

Description

The z/OS LDAP server, when using the extended operation (EXOP) backend, supports two LDAP extended operations, GetDnForUserid and GetPrivileges, that retrieve Policy Director data from any LDAP server.

What this change affects

You can retrieve Policy Director data. Applications on z/OS can use the AZN APIs and retrieve Policy Director data using the LDAP extended operations.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

None.

For more information

See Chapter 22, "Using extended operations to access Policy Director data" on page 247 and Appendix H, "Supported extended operations" on page 435.

APF authorization

Description

The z/OS LDAP server PDSs must be APF-authorized prior to running the LDAP server.

What this change affects

With the proper APF-authorizations, the LDAP server can make the necessary threading calls.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

APF-authorize the LDAP server PDS and any other PDSs or DLLs that get invoked from the LDAP server.

For more information

See "Protecting the environment for the LDAP server" on page 37.

listen configuration option

Description

The z/OS LDAP server provides a manner of binding and listening on multiple IP addresses and ports with the **listen** configuration option. This migration is optional, but highly recommended.

What this change affects

Customization of your configuration file.

Dependencies

Must have APAR OW50971 (LDAP) for PC call support.

Coexistence considerations

None.

Migration tasks

Removal of port, securePort, and security from the configuration file and replacing with one or more equivalent listen parameters.

For more information

See "Migrating from port, securePort, and security to listen option" on page 106 and the listen option on page 59.

commThreads migration

Description

The maxThreads and waitingThreads options have been deprecated and are no longer evaluated by the LDAP server because of the new threading model in the LDAP server. The commThreads parameter defines the number of threads that the LDAP server initializes for communicating with clients.

What this change affects

Customization of your configuration file.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

Removal of maxThreads and waitingThreads from the configuration file and replacing it with commThreads.

For more information

See "Migrating from maxThreads and waitingThreads to commThreads option" on page 107 and the commThreads option on page 55.

maxConnections migration

Description

The maxConnections option has been updated to no longer be dependent upon maxThreads and reflects the expected maximum number of concurrently connected clients.

What this change affects

Customization of your configuration file.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

Removal of maxThreads from the configuration file and updating the maxConnections number.

For more information

See "Migrating to the updated maxConnections option" on page 108 and the maxConnections description on page 62.

z/OS V1R1 overview

No new function was introduced to the z/OS LDAP server for this release.

OS/390 V2R10 overview

This section describes the new and changed LDAP server functions that were introduced for OS/390 V2R10 (V2R10). The information about each item includes:

- Description
- Summary of the LDAP server tasks or interfaces that may be affected
- · Coexistence considerations, if any, that are associated with the item
- · Migration procedures, if any, that are associated with the item
- · References to other publications that contain additional detailed information.

TDBM backend

Description

The LDAP server comes with a TDBM backend database based on DB2. The TDBM database is a highly scalable database implementation.

What this change affects

The amount of data the directory can store.

Dependencies

None.

Coexistence considerations

None.

Migration tasks

Unload RDBM using db2ldif on the release migrating from. Load resulting LDIF file into TDBM using Idif2tdbm or Idapadd on the release migrating to. TDBM must be configured. The schema must be determined.

For more information

See "Installing and setting up DB2 for TDBM" on page 13, "Protecting the environment for the LDAP server" on page 37, "Creating the DB2 database and table spaces for TDBM" on page 40, Chapter 11, "Running and using the LDAP backend utilities" on page 125, and Chapter 16, "LDAP directory schema for TDBM" on page 159.

Schema publication

Description

The LDAP server allows the schema to be retrieved using the LDAP protocol. Searches can be directed to both the TDBM backend and the SDBM backend.

What this change affects

LDAP clients can learn the schema being used by the LDAP server.

Dependencies

TDBM or SDBM must be configured.

Coexistence considerations

None.

Migration tasks

None.

For more information

See "Displaying the schema entry" on page 175.

Schema update

Description

The LDAP server allows the schema to be changed dynamically through the LDAP protocol when using a TDBM backend.

What this change affects

Reduces server outages by permitting dynamic schema update.

Dependencies

TDBM must be configured.

Coexistence considerations

None.

Migration tasks

Determine if the schema in use matches the IBM-shipped schema. If so, load the IBM-shipped schema in the new format. If additions or changes have been made, migrate the schema using the schema2ldif utility.

For more information

See Chapter 16, "LDAP directory schema for TDBM" on page 159 and "Migrating the schema from RDBM to TDBM" on page 176.

LDAP server configuration utility

Description

The LDAP configuration utility, **Idaponf**, simplifies and automates the LDAP server configuration process for TDBM and, optionally, SDBM.

What this change affects

With this method you can configure the LDAP server with minimal user interaction and allow novice LDAP users to quickly deploy an LDAP server.

Dependencies

The **Idapenf** utility has the following dependencies:

- · Generates a procedure; therefore, the LDAP server must run as a started task.
- Assumes that RACF is the security server in use.
- · Does not handle multiple TDBM backends.
- All values in the input files must be less than 66 bytes in length and must contain only printable characters in the IBM-1047 code page.
- Cannot extend or enhance an existing LDAP server configuration. Furthermore, any manual updates to the output that the utility produces will be lost if you run the utility again with the same output data set.

Convenience considerations

None.

Migration tasks

None.

For more information

See Chapter 4, "Configuring an LDAP server using the Idapcnf utility" on page 17.

Summary of interface changes

This section summarizes the new and changed interface components of the LDAP server. Such as:

- Commands
- Utilities
- · Configuration options

nterface change:	Release:	Refer to page:
digestRealm, logfile,, and serverEtherAddr configuration options new)	z/OS V1R4	53
attribute, index, objectclass, tbspaceentry, tbspacemutex, bspace32k, tbspace4k, and verifySchema configuration options deprecated)	z/OS V1R4	
dapcp command (removed)	z/OS V1R4	
dapadduuids utility (new)	z/OS V1R4	141
dbm2ldif utility (changed)	z/OS V1R4	
db2ldif and ldif2db utilities (removed)	z/OS V1R4	
sslCipherSpecs (changed)	z/OS V1R4	
Message severities (changed)	z/OS V1R4	
Return codes (changed)	z/OS V1R4	297
krbKeytab, krbLDAPAdmin, listen, nativeAuthSubtree, nativeUpdateAllowed, pcThreads, serverKrbPrinc, supportKrb5,		
and useNativeAuth configuration options (new)	(00.11.5)	
port, securePort, and security configuration options (deprecated)	z/OS V1R2	73
port, securePort, and security configuration options (deprecated) maxThreads and waitingThreads (ignored)	z/OS V1R2	73
port, securePort, and security configuration options (deprecated)		
port, securePort, and security configuration options (deprecated) maxThreads and waitingThreads (ignored)	z/OS V1R2	73
maxThreads and waitingThreads (ignored) slapd -I command-line parameter (new)	z/OS V1R2 z/OS V1R2	73 94
maxThreads and waitingThreads (ignored) slapd -I command-line parameter (new) slapd -d command-line parameter (changed)	z/OS V1R2 z/OS V1R2 z/OS V1R2	73 94 94
maxThreads and waitingThreads (ignored) slapd -I command-line parameter (new) slapd -c command-line parameter (changed) slapd -s, -p command-line parameters (deprecated)	z/OS V1R2 z/OS V1R2 z/OS V1R2 z/OS V1R2	73 94 94 94
maxThreads and waitingThreads (ignored) slapd -I command-line parameter (new) slapd -d command-line parameter (changed) slapd -s, -p command-line parameters (deprecated) dapcnf configuration utility (new)	z/OS V1R2 z/OS V1R2 z/OS V1R2 z/OS V1R2 OS/390 V2R10	73 94 94 94 17
coort, securePort, and security configuration options (deprecated) maxThreads and waitingThreads (ignored) slapd -I command-line parameter (new) slapd -d command-line parameter (changed) slapd -s, -p command-line parameters (deprecated) dapcnf configuration utility (new) ettrOverflowSize configuration option (new)	z/OS V1R2 z/OS V1R2 z/OS V1R2 z/OS V1R2 OS/390 V2R10 OS/390 V2R10	73 94 94 94 17 53
coort, securePort, and security configuration options (deprecated) maxThreads and waitingThreads (ignored) slapd -I command-line parameter (new) slapd -d command-line parameter (changed) slapd -s, -p command-line parameters (deprecated) dapcnf configuration utility (new) attrOverflowSize configuration option (new) sslCertificate configuration option (new)	z/OS V1R2 z/OS V1R2 z/OS V1R2 z/OS V1R2 OS/390 V2R10 OS/390 V2R10 OS/390 V2R10	73 94 94 94 17 53 53

Chapter 11. Running and using the LDAP backend utilities

Utility programs are provided to assist in initializing and backing up the data managed by the LDAP server.

Operation	TDBM utility
Load data into backend database	ldif2tdbm
Unload data from backend database to an LDIF file	tdbm2ldif

These programs can be run in the z/OS shell, as jobs using JCL and procedures, or from TSO.

Format and usage information for the utilities are in:

- "ldif2tdbm program" on page 127
- "tdbm2ldif program" on page 137

See "Creating the DB2 database and table spaces for TDBM" on page 40 for information on the permissions necessary to run these utilities.

Running the LDAP TDBM backend utilities in the z/OS shell

In order to run the **Idif2tdbm** or **tdbm2Idif** utilities in the shell, some environment variables need to be set properly. Sample shell scripts are shipped with the LDAP server for **Idif2tdbm** and **tdbm2Idif**. The sample shell scripts are located in **/usr/lpp/Idap/sbin**. These shell scripts need to be modified to fit your environment. The **PATH** and **STEPLIB** variables should be set in the shell scripts. Ensure that **/usr/sbin** is added to the **PATH** environment variable. Also, be sure **STEPLIB** is set to **GLDHLQ**.SGLDLNK if the PDS has not been placed in LINKLIST.

If you are using the **Idif2tdbm** utility to add an entry containing a password that will be encrypted, the **LIBPATH** variable may need to be set. See "Installing OCSF and ICSF for password encryption" on page 16 for more information.

When started, these utilities read an environment variable file. The default file is /etc/ldap/slapd.envvars. This default can be changed by setting the environment variable LDAP_SLAPD_ENVVARS_FILE to the full path name of the desired environment variable file.

Running the LDAP TDBM backend utilities from JCL

Sample JCL for running **ldif2tdbm** and **tdbm2ldif** from batch is provided with the LDAP server. The JCL includes an inline procedure, which will need to be modified by each installation to ensure that the **ldif2tdbm** and **tdbm2ldif** load modules can be found. It may also be necessary to modify the **JOB** card for installation-specific requirements. Appendix E, "Sample JCL" on page 413 shows each sample JCL, which includes instructions for passing parameters into the utility programs. These jobs can be run by editing the JCL member of the *GLDHLQ*.SGLDSAMP dataset and entering the **submit** command.

Notes:

 If your runtime libraries for DB2 are not in LINKLIB or LPA on the system, make sure you specify the DB2 high-level qualifier for your DB2 installation in a STEPLIB DD card in the LDF2TDBM or TDBM2LDF batch job. The LDAP server and utilities require the following DB2 dataset: DB2HL0.SDSNLOAD

Running the LDAP TDBM backend utilities in TSO

The Idif2tdbm and tdbm2ldif utilities can be run from TSO. Following are the steps to do this:

- 1. Specify the PDS (GLDHLQ.SGLDLNK) where the LDAP server load modules are installed: tsolib act dsn('GLDHLQ.SGLDLNK')
- 2. Make the PDS (GLDHLQ.SGLDEXEC) containing the CLISTs needed to run the utilities available in SYSEXEC:

```
alloc f(SYSEXEC) da('GLDHLQ.SGLDEXEC')
```

3. You can change the default environment variable file for the utilities by creating a dataset to hold the environment variables and then using the TSO alloc command as shown:

```
alloc da(ENVAR) dsn('datasetname')
```

If you want to specify a configuration file that is a data set you can specify the -f option on the command. For example:

```
tdbm2ldf -f "//'datasetname'" -o /tmp/ldif.1
```

Alternately, to specify the configuration file by associating it with a DD name, enter in TSO:

```
alloc da('datasetname') fi(config) shr
```

and then invoke the utility without specifying the -f option.

Once this setup is complete, running these utilities follows the same syntax as would be used if running in the z/OS shell, except you must use Idf2tdbm instead of Idif2tdbm and tdbm2Idf instead of tdbm2Idif. See "Running the LDAP TDBM backend utilities in the z/OS shell" on page 125.

Idif2tdbm program

Purpose

This program is used to load entries specified in LDAP Data Interchange Format (LDIF) into a TDBM directory stored in a relational database. The database must already exist. The **Idif2tdbm** program is intended for loading a large number of entries. The utility creates load records from the entries in the LDIF input files, then runs the DB2 Load Utility to load the records into the TDBM database.

The Idapadd command can also be used to add entries to a TDBM database. See "When to use ldif2tdbm" on page 131 for more information on when to use Idif2tdbm or Idapadd. See z/OS: Security Server LDAP Client Programming for more information on Idapadd.

The **Idif2tdbm** program may be used to add entries to an empty directory database, or to a database that already contains entries. The Idif2tdbm utility may also be used to modify the schema entry before loading new entries into the database.

See "Preparing to run Idif2tdbm" on page 129 before using Idif2tdbm.

Format

```
ldif2tdbm {[-c [-n noe|nop|noep]] [-p] [-1]}
-o outHlq {[-i ldifFile[,ldifFile]...] [-e ldifListFile]}
[-s schemaLdifFile], schemaLdifFile]...] [-v schemaListFile] [-f confFile] [-a yes|no] [-t logFile]
[-b creatorDN] [-q summaryFrequency] [-d debugLevel]
```

Parameters

-c Check that the entries in the LDIF input file or files are complete and acceptable according to the current internal schema. If the -s option is specified, the current internal schema is updated using the entries in the schema LDIF files before the LDIF entries are checked.

-n noe | nop | noep

Control the amount of checking that is done for each LDIF entry. In some cases, the checking that is skipped is performed anyway during the prepare step. See the specific information after each value.

- noe Do not perform entry existence checks:
 - Do not check that the entry does not already exist, either as a previous LDIF entry or in the LDAP database.

Note that a previous duplicate LDIF entry will be detected and rejected during the prepare step (when -p is specified). However, if **noe** is specified, an existing entry in the database will not be detected. Thus the load step (when -I is specified) can result in a duplicate entry in the LDAP database.

- nop Do not perform several parent checks:
 - Do not check that the entry's parent exists, either as a previous LDIF entry or in the LDAP database.
 - Do not check that the LDIF entry is not under a referral entry.

Note that the checks skipped by **nop** are always performed during the prepare step (when **-p** is specified).

- noep Do not perform these checks:
 - Do not check that the LDIF entry already exists.
 - Do not check that the entry's parent exists.
 - Do not check that the entry is not under a referral entry.
- -p Prepare DB2 table load files and JCL from the entries in the LDIF input file or files. The load and JCL files are generated as datasets, whose high-level qualifier is specified with the **-o** option. The prepare

ldif2tdbm

step deletes the existing contents of the datasets before writing new contents to the datasets. If the -s option is specified, the current internal schema is updated using the entries in the schema LDIF files before the LDIF entries are prepared.

-I Invoke the DB2 Load Utility to load the LDAP database. Also, update the schema entry in the database if -s is specified.

-o outHla

Specify the high level qualifier of the datasets that will contain the DB2 load and JCL output, the status information, and the system information. These datasets must be allocated before invoking Idif2tdbm. See "Preparing to run Idif2tdbm" on page 129 for more information.

-i IdifFile

Specify the name of an LDIF file to use as input. If a list of file names is specified, each file in the list is processed in turn. If **Idif2tdbm** is invoked separately to run each **Idif2tdbm** step (check, prepare, and load), make sure to specify the same set of LDIF files each time you run Idif2tdbm. The **Idif2tdbm** program issues a warning prompt (unless -a is specified) if the list of LDIF file names is different.

See "Using LDIF format to represent LDAP entries" on page 290 for more information on the general format of LDIF entries.

-e IdifListFile

Specify the name of a file which contains a list of LDIF input files to be used as input. This is equivalent to using the -i option to specify a list of LDIF files, but is more convenient for a large list. Each record in the list file must contain the name of one LDIF input file. Blank lines and lines beginning with a # (comment lines) are ignored. See the -i explanation above for more information about using a list of LDIF files.

-s schemaLdifFile

Specify the name of an LDIF file containing only schema entries. The entries must be in **Idapmodify** LDIF-mode format (see Note 8 on page 135). If a list of file names is specified, each file in the list is processed in turn. During the check or prepare steps (when -c or -p is specified), the current internal schema is updated using these schema files, but the schema is not changed in the database. This updated internal schema is used to check the entries in the LDIF input files or to prepare the load files. During the load step (when -I is specified), these schema files are used to update the current schema in the database. If Idif2tdbm is invoked separately to run each Idif2tdbm step (check, prepare, and load), make sure to specify the same set of LDIF schema files each time you run Idif2tdbm. The Idif2tdbm program issues a warning prompt (unless -a is specified) if the list of LDIF file names is different.

-v schemaListFile

Specify the name of a file which contains a list of LDIF input files to be used to modify the schema entry. This is equivalent to using the -s option to specify a list of LDIF files, but is more convenient for a large list. Each record in the list file must contain the name of one LDIF input file. Blank lines and lines beginning with a # (comment lines) are ignored. See the -s explanation above for more information about using a list of LDIF files.

-f confFile

Specify the name of the LDAP configuration file to use. This configuration file only needs to contain information for the TDBM backend into which the entries are to be loaded. The default is /etc/ldap/slapd.conf.

-a yes|no

Eliminate the prompt that **Idif2tdbm** issues when it detects an unexpected status condition by providing an answer to the prompt. If an unexpected status condition is encountered, Idif2tdbm continues if -a yes is specified and exits if -a no is specified. If -a is not specified, Idif2tdbm issues a prompt and waits for a response. See Note 11 on page 135 for more information on status checking.

-t logFile

Specify the name of the file to which messages are written. If -t is not specified, messages are sent to stdout or stderr. The existing contents of the file are deleted.

-b creatorDN

The DN to be associated as creator with each loaded entry that does not already include a creatorsname attribute. If -b is not specified, the value of the adminDN option in the LDAP server configuration file is used. If neither option is specified, Idif2tdbm fails. The same processing is also applied for the modifier for each loaded entry that does not already include a modifiersname attribute.

-q summaryFrequency

Specify the number of LDIF entries the prepare step should process between issuing summary messages. The default value is 1000, resulting in issuing a summary message after every 1000 entries are prepared. If you have many LDIF entries, increase this value to reduce the frequency of summary messages. Specify a negative value or 0 to suppress issuing any intermediate summary messages. A final summary message is always issued.

-d debugLevel

Specify the level of debug messages to be created. The debug level is specified in the same fashion as the debug level for the LDAP server, as described on page 94. Table 16 on page 95 lists the specific debug levels. The default is no debug messages.

All other command line inputs will result in a syntax error message, after which the proper syntax will be displayed. Also, specifying the same option multiple times with different values results in a syntax error message.

Examples

Following is an example using the **ldif2tdbm** utility:

ldif2tdbm -cpl -i /data2/ldif.data -s /data2/schema changes.data -o admin3.prv -d ERROR

This Idif2tdbm utility invocation checks, prepares, and loads the LDIF data from /data2/1dif.data, updates the schema using the LDIF data from /data2/schema changes.data, uses the output datasets ADMIN3.PRV.BULKLOAD.INPUT.xxx and ADMIN3.PRV.BULKLOAD.JCL, and specifies a debug level of ERROR for LDAP_DEBUG_ERROR. By default, the configuration file used is /etc/ldap/slapd.conf, complete checking is done on the entries, the user is prompted whether to continue if warning conditions are detected, messages are sent to stdout or stderr, the value of the adminDN option in the configuration file is used as the creator of each loaded entry, and a progress message is issued after every 1000 entries processed.

Preparing to run Idif2tdbm

Before invoking **Idif2tdbm**, you must

- · Allocate the output datasets used by Idif2tdbm.
- Create the SYSTEM file used by the prepare step of **Idif2tdbm**.
- Set the environment variables used by **Idif2tdbm**.

Allocating datasets required by Idif2tdbm

The various Idif2tdbm processing steps use 5 load record datasets and a JCL dataset. These datasets must be allocated before invoking Idif2tdbm and the high-level qualifiers in the dataset names must be specified as the value of the -o outHIq option on the command. The Idif2tdbm utility writes the contents of these datasets, except for the SYSTEM member of the JCL dataset which must be created before Idif2tdbm is run.

· Load Record Datasets

The prepare step of Idif2tdbm writes the load data created from each LDIF entry into the load record datasets. Each dataset contains the records for one database table and is used as input to the DB2 Load Utility when loading that table.

Dataset names:

outHlq.BULKLOAD.INPUT.DESC (DIR_DESC load dataset)

Idif2tdbm

outHlq.BULKLOAD.INPUT.ENTRY (DIR_ENTRY load dataset) outHig.BULKLOAD.INPUT.LATTR (DIR LONGATTR load dataset) outHIq.BULKLOAD.INPUT.LENTRY (DIR_LONGENTRY load dataset) outHiq.BULKLOAD.INPUT.SEARCH (DIR_SEARCH load dataset)

Dataset format:

Sequential (non-PDS) Record format = VB

Dataset record length and block size:

The record length depends on the page size of the corresponding table space. The actual record length may be reduced slightly by Idif2tdbm at run-time.

For a 32K page size in DB2, use LRECL=32756, BLKSIZE=32760.

This record length and block size will also work for smaller page sizes. However, for smaller page sizes, a smaller LRECL and BLKSIZE can be used to reduce the required disk space. For example, on 3390 DASD, LRECL=27994, BLKSIZE=27998 will allow 2 blocks per track and, in general, more bytes per track to be written. For the smaller page sizes, the LRECL and BLKSIZE must be at least pagesize + 6 and pagesize + 10, respectively.

Dataset size:

You may use the space estimation tool to determine the approximate size of the load data sets. This tool can be downloaded by selecting "Download" on:

http://www.ibm.com/software/network/directory

The tool is written in awk script and can be run under z/OS UNIX System Services. The script can be modified for your specific needs based on values specific to your configuration.

Alternatively, a rough estimate for the size (in bytes) of each dataset is as follows:

```
outHLQ.BULKLOAD.INPUT.DESC:
```

34 * (average depth) * (number of entries in the LDIF files)

where average depth = (average number of levels in a DN) - (average number of levels in each DN's suffix) + 1

outHLQ.BULKLOAD.INPUT.ENTRY, outHLQ.BULKLOAD.INPUT.LATTR, and outHLQ.BULKLOAD.INPUT.LENTRY:

The combined space required for these files is roughly

(number of bytes in the LDIF input files) * 3.0

If most directory entries are shorter than the DIR ENTRY page size, and most attributes are shorter than the attrOverflowSize in the LDAP server configuration file, then most of the data will be written in the outHLQ.BULKLOAD.INPUT.ENTRY dataset. Otherwise, you may wish to allocate each of these three datasets as this maximum size to ensure that you do not run out of space.

outHLQ.BULKLOAD.INPUT.SEARCH:

(number of bytes in the LDIF input files) * 2.5

· JCL dataset

The JCL dataset is a PDS whose members contain system information, status information, and the JCL to run DB2 Load Utility for each database table that needs to be loaded. All steps of Idif2tdbm use the status information. The prepare step writes the contents of the JCL members of the JCL dataset, using the information in the system member. The system member of the JCL dataset must be created by the user before **Idif2tdbm** is invoked to run the prepare step.

Dataset name:

outHlq.BULKLOAD.JCL

Members:

outHlq.BULKLOAD.JCL(JDESC) (DIR_DESC load JCL) outHIq.BULKLOAD.JCL(JENTR) (DIR_ENTRY load JCL) outHlq.BULKLOAD.JCL(JLATT) (DIR_LONGATTR load JCL) outHlq.BULKLOAD.JCL(JLENT) (DIR LONGENTRY load JCL) outHIq.BULKLOAD.JCL(JSRCH) (DIR_SEARCH load JCL) outHlq.BULKLOAD.JCL(STATUS) (Status information) outHlq.BULKLOAD.JCL(SYSTEM) (System information - see below)

Dataset format:

Partitioned (PDS) Record format = FB Record length = 80

Dataset size:

200K bytes

Creating the SYSTEM file

The contents of the SYSTEM file, outHlq.BULKLOAD.JCL(SYSTEM), must be created before invoking Idif2tdbm to run the prepare step, which uses the information in the SYSTEM file to create the JCL to invoke the DB2 Load Utility to load each of the database tables.

· Format of SYSTEM records

SSID yyy Subsystem ID for DB2

HLQ db2hlq High level qualifier for DB2 datasets

JOBCARD //jobname... Job card record for JCL

Ignored if first non-blank character is # #Comment record

- The SYSTEM member must contain one SSID record, one HLQ record, and one or more JOBCARD records. The first JOBCARD record must begin with //jobname where jobname is at most 8 characters. The maximum length for each record value is 55 for the SSID value, 36 for the HLQ value, and 71 for each JOBCARD value. Make sure there are no sequence numbers at the end of each line.
- To differentiate the load jobs in the 5 JCL members of the JCL PDS, Idif2tdbm replaces the last character of the job name with the digits 1 through 5, as follows:

```
Jobname ends in 1 - JCL for loading DIR_DESC table, in outHlq.BULKLOAD.JCL(JDESC)
Jobname ends in 2 - JCL for loading DIR_ENTRY table, in outHlq.BULKLOAD.JCL(JENTRY)
Jobname ends in 3 - JCL for loading DIR_LONGATTR table, in outHlq.BULKLOAD.JCL(JLATT)
Jobname ends in 4 - JCL for loading DIR_LONGENTRY table, in outHlq.BULKLOAD.JCL(JLENT)
Jobname ends in 5 - JCL for loading DIR_SEARCH table, in outHlq.BULKLOAD.JCL(JSRCH)
```

Setting environment variables used by Idif2tdbm

The Idif2tdbm utility uses two environment variables. These environment variables are optional, but if set properly, they may increase performance.

LDIF2TDBM CACHE SIZE

The LDIF2TDBM CACHE SIZE environment variable sets the number of internal structures that **Idif2tdbm** should allocate. This number should be set to two times the number of entries that will be loaded.

LDIF2TDBM LOG DB

The LDIF2TDBM LOG DB environment variable determines whether Idif2tdbm will force the DB2 Load Utility to do logging as it loads the database. See Note 5 on page 133 for more information.

When to use Idif2tdbm

Both the **Idif2tdbm** utility and the **Idapadd** command can be used to load entries into a TDBM database. There are advantages and disadvantages to each of these. Following is a list of considerations that can assist in determining which method to use when loading entries.

Idif2tdbm

Table 20. Considerations for using Idif2tdbm or Idapadd

Considerations	ldif2tdbm	ldapadd
Speed of load	Faster	Slower, especially if the database is already large.
Operational attributes	Accepts input containing creatorsname , modifiersname , and ibm-entryuuid operational attributes.	Input cannot contain these operational attributes.
Complexity	High, will take time to learn. Normal usage requires multiple invocations, with review of JCL and preparation for recovery before load.	Low
Set up	User must allocate datasets and create the SYSTEM information file before running Idif2tdbm for the first time.	No set up.
LDAP server down time	LDAP server must be down during load step. Server can be up during check and prep steps.	Server must be operational during adds.
DB2 logging	Logging is optional. Additional DB2 work is needed to make database fully usable after the load if logging is not done.	Requires logging.
Recovery	Recovery from load failure is complex, involving knowledge of DB2 Utilities.	Simple recovery.
Invocation	Must be run from z/OS system containing the database, or any image in a Parallel Sysplex running a DB2 subsystem which is a member of the DB2 data sharing group containing the database, using a user ID with DB2 privileges.	Can be run from any LDAP client with appropriate LDAP access.
Re-usability	Prepared entries are saved, so they can be re-used to load another system or reload this system.	No saved output.
Replication	New entries are not replicated.	New entries can be replicated.

Summary: The Idif2tdbm utility usage is complex, but it is fast. The Idapadd utility usage is simpler, but it is slower.

Recommendation: For one-time additions of 100K or more entries, or for frequent additions of 10K or more entries, use Idif2tdbm. For infrequent additions of less than 10K entries, use Idapadd. For additions of between 10K and 100K entries, use either Idif2tdbm or Idapadd.

Idif2tdbm performance considerations

The **Idif2tdbm** utility performance considerations are:

- 1. Idif2tdbm uses a lot of storage when adding a large number of entries. Make sure you have sufficient memory available.
- 2. If the parent of an entry being added is in the TDBM database, then Idif2tdbm must also check the database to ensure that the entry itself does not already exist. Thus, to minimize the database checks, include the parents of the entries being added in the **Idif2tdbm** input whenever possible. For example, do not use Idapadd to add a suffix and then use Idif2tdbm to add all the entries under the suffix. Instead, include the suffix in the Idif2tdbm input.
- 3. Do not specify the debug command line option, -d. Usage of debug can impact performance even when there is little debug output. Only specify -d when an error has occurred that you cannot fix.
- 4. If you know that your LDIF input is acceptable, consider using the -n option to limit the amount of checking done during the check (-c) step. You can also combine the check (-c) and prepare (-p) steps rather than doing them separately. In particular, using the -cp -n noe options will reduce the checks that access the database. This is especially useful if the parents of the entries being added are not in the **Idif2tdbm** input (see 2 above).

5. By default, for better load performance Idif2tdbm sets the LOG NO option in the DB2 Load Utility JCL created during the prepare (-p) step. When the load (-I) step invokes the DB2 Load Utility, the Load Utility will set the copy pending restriction against the table space. The database can be read but cannot be updated until the restriction is removed, for example, by running the DB2 REORG or COPY utility. To avoid entering the copy pending state (for example, when the database already contains a large number of entries), change LOG NO to LOG YES in the JCL. This will be done by Idif2tdbm if the LDIF2TDBM_LOG_DB environment variable is set to 1 before running the prepare step. It can also be done manually in each of the outHLQ.BULKLOAD.JCL members after the prepare step if the load step is run separately from the prepare step.

Idif2tdbm normal usage

The normal usage of **Idif2tdbm** is:

- 1. Perform the necessary setup for running Idif2tdbm, as described in the "Preparing to run Idif2tdbm" on page 129. This consists of:
 - allocating datasets required by Idif2tdbm
 - creating the SYSTEM member in the outHLQ.BULKLOAD.JCL dataset
 - exporting the LDIF2TDBM_CACHE_SIZE and LDIF2TDBM_LOG_DB environment variables, if desired.

Note that the same datasets can be used for loading different entries, but the contents of the datasets will be overwritten each time.

- 2. Repeatedly invoke **Idif2tdbm** with only the check (-c) option until all problems in the LDIF input files have been resolved. The check step can be skipped if you are sure that the LDIF input files contain only valid entries; however, although the prepare step does some checking of each entry, it might not detect all problems and might allow an entry that is not valid to be added to the database.
- 3. Run **Idif2tdbm** with the prepare (**-p**) option to prepare the load data.
- 4. Bring the LDAP server down.
- 5. Even if loading an empty database, make a full image copy of the table spaces for the DIR DESC, DIR SEARCH, DIR ENTRY, DIR LENTRY, and the DIR LATTR tables. A full image copy is required to help recover from any potential DB2 Load Utility failures. It is done after the Idif2tdbm invocation with the prepare (-p) option specified because the prepare (-p) option may update some of the tables Idif2tdbm is attempting to load. Thus, these updates must be captured for a successful recovery from a DB2 Load Utility failure. If the prepare (-p) and load (-I) options are specified together, the full image copy should be done before the **Idif2tdbm** invocation.
- 6. Review the JCL created by the prepare step in the JCL dataset, outHLQ.BULKLOAD.JCL. Ensure that the DB2 Load Utility work datasets are large enough and that the DB2 Load Utility options, especially the LOG value, are acceptable. See note 5 in the Idif2tdbm performance considerations section for more information on the LOG value.
- 7. Run **Idif2tdbm** once more, with the load (-I) option to load the data into the database. This will submit 5 batch jobs to load the database.
- 8. Review the output from the load jobs when they terminate. If the loaded table spaces are in the copy pending state, refer to DB2 for OS/390 and z/OS: Utility Guide and Reference for instructions on removing that restriction on table spaces. This usually involves running a DB2 utility to create an image copy of the database or reorganize the database (or both).
- 9. Run the DB2 runstats utility to reset the statistics used by DB2 to access the database.
- 10. Run the DB2 copy utility to make an backup image copy of the database.

Idif2tdbm recovery

The Idif2tdbm program can determine whether the submission of the DB2 Load Utility jobs was successful, but it cannot determine if the DB2 Load Utility jobs themselves succeed. In fact, Idif2tdbm normally terminates before the jobs are finished. Thus, the final **ldif2tdbm** success message will appear even if one or more of the jobs eventually fails.

Idif2tdbm

If a DB2 Load Utility job fails, terminate all of the DB2 Load Utility invocations that failed, using -TERM UTIL(BULKx)

where x is the last number in the job name of the failing job. Then, there are two alternatives for recovery.

Recovery process 1

The first recovery alternative is not selective in nature and may result in unnecessary reloading of data. This alternative must be used for recovery when:

- The **Idif2tdbm** utility is run with the load (**-I**) option immediately after creating the TDBM database (that is, without first starting the server or running **Idif2tdbm** with the check (**-c**) or prepare (**-p**) options).
- Loading the DIR_ENTRY, DIR_LENTRY, or DIR_LATTR table fails and a full image copy of these tables
 does not exist.

Following is the first recovery alternative process:

- 1. If full image copies exist for all 5 table spaces, use the DB2 Recover Utility to recover all 5 table spaces from those full image copies. Otherwise, drop and recreate the DIR_ENTRY, DIR_LATTR, DIR_SEARCH, DIR_DESC, and DIR_LENTRY tables.
- 2. Refer to *DB2 for OS/390 and z/OS: Utility Guide and Reference* for information on correcting the problems logged in the failing jobs.
- 3. Run Idif2tdbm with only the load (-I) option again.

Recovery process 2

The second recovery alternative requires less processing, but requires a greater understanding of the problem. It cannot be used for recovery when:

- The Idif2tdbm utility is invoked with the load (-I) option immediately after creating the TDBM database.
- Loading the DIR_ENTRY, DIR_LENTRY, or DIR_LATTR table fails and a full image copy of these tables
 does not exist.

Following is the second recovery alternative process:

- If the load jobs that failed involve only the DIR_SEARCH or DIR_DESC tables, follow these steps:
 - 1. If a full image copy exists for the failing table spaces, use the DB2 Recover Utility to recover these table spaces from those full image copies. Otherwise, drop and recreate the failing tables.
 - 2. Refer to *DB2 for OS/390 and z/OS: Utility Guide and Reference* for information on correcting problems logged in the failing jobs.
 - 3. Manually resubmit the DB2 Load Utility jobs that failed.
 - 4. Make sure that the DIR_DESC table contains a row with the values (-2, -2).
- If the load jobs that failed involve the DIR_ENTRY, DIR_LENTRY, or DIR_LATTR tables and a full
 image copy exists (you must use the first recovery alternative if a full image copy does not exist), follow
 these steps:
 - 1. If a schema file was specified using the **-s** option, the DIR_ENTRY, DIR_LENTRY, and DIR_LATTR tables must **all** be recovered, along with other failing tables. If a schema file was not specified, only the failing tables need to be recovered.
 - 2. If a full image copy exists for the affected table spaces, use the DB2 Recover Utility to recover these table spaces from those full image copies. Otherwise, you must use the first alternative.
 - 3. Refer to *DB2 for OS/390 and z/OS: Utility Guide and Reference* for information on correcting problems logged in the failing jobs.
 - 4. Manually resubmit the DB2 Load Utility jobs for all the affected tables.
 - 5. Make sure that the DIR_DESC table contains a row with the values (-2, -2).
 - 6. If a schema file was specified using the **-s** option, follow these steps:

- a. If the table spaces are in the copy pending state, refer to DB2 for OS/390 and z/OS: Utility Guide and Reference for information on removing that restriction on the table spaces.
- b. Start the LDAP server.
- c. Remove the 'version: 1' record from the schema file, if it contains one. Then, perform an LDAP modify using the schema file as input.

Usage

- 1. All input files specified with the -i, -e, -s, and -v options can be HFS files or datasets. In order to use datasets with Idif2tdbm, you must use a VB record format. Further, separator lines between directory entries in the LDIF file must contain only space characters (X'40'), if anything.
- 2. The Idif2tdbm utility can be invoked to run the check (-c) and prepare (-p) steps while an LDAP server using the same TDBM database is active. The LDAP server must be down when the load (-I) step is run.
- 3. The LDAP_DEBUG environment variable can also be used to set the debug level for Idif2tdbm. See page 94 for more information on specifying the debug level.
- 4. The DB2 runstats utility should be run once data is loaded so that DB2 queries are optimized.
- 5. No replication is performed for entries added by **Idif2tdbm**. A new replica entry can be added using **Idif2tdbm**, but replication does not begin until the LDAP server is started.
- 6. Referral entries can be added by **Idif2tdbm**.
- 7. Only one TDBM database is loaded in an invocation of Idif2tdbm. All the LDIF entries in the LDIF input files and schema input file must contain DNs that belong to the same TDBM database. The parent of each LDIF entry must either be already in the database or must be a prior LDIF entry within these LDIF input files.
- 8. Restrictions for updating the schema directory entry using **Idif2tdbm**:
 - The schema can only be updated within an Idif2tdbm invocation by using the -s schemaLdifFile and -v schemaListFile options. The LDIF input files specified in the -i and -e options cannot contain a schema entry.
 - The LDIF schema files specified using the -s and -v options must contain only schema entries, in the LDIF mode format supported by Idapmodify, described in z/OS: Security Server LDAP Client *Programming* with the following changes:
 - If the optional **changetype** line is specified, the value must be **modify**. No other **changetype** value is supported for the schema.
 - The change indicator line (add:x, replace:x, or delete:x) is not optional and there is no default change indicator. Each change_clause must begin with a change indicator line.
- 9. The aclSource and ownerSource attributes should not be specified in an LDIF entry and are ignored. These attributes are set only by the system.
- 10. The Idif2tdbm utility check processing is terminated after 100 syntax errors are detected. The Idif2tdbm prepare and load processing are terminated after the first error. These values cannot be modified.
- 11. Since normal Idif2tdbm usage can involve multiple invocations to check, prepare, and load the entries, Idif2tdbm maintains a status file to keep information about the last successful step processed. The status file is outHlq.BULKLOAD.JCL(STATUS), where outHlq is the value of the -o option on the command. Any invocation of Idif2tdbm with the same value for -o is considered a continuation of processing of an earlier invocation.
 - Before processing any entries, Idif2tdbm uses the information about the earlier invocation in the status file to determine that each processing step is performed in order, the input files used for this invocation are the same as the ones used for the earlier invocation, and this invocation is not going to delete output prepared by the earlier invocation.
 - The Idif2tdbm utility issues a warning message for each condition that it detects, followed by a single prompt asking if processing should continue. The prompt can be suppressed by specifying the **-a yes|no** option on the command to provide an answer for the prompt.

Idif2tdbm

- · At the end of processing, Idif2tdbm rewrites the status file with information about the last successful step processed (check, prepare, or load), except as follows:
 - If the check step succeeds but
 - Only the check (-c) step is requested, or the check (-c) and load (-l) steps are requested and the load step fails
 - The previous status is prepare (P) or load (L)

the status file is not rewritten, to avoid replacing a higher status with a lower one.

- If no step succeeds but
 - The prepare (-p) step fails
 - The previous status is prepare (P) or load (L)

the status in the status file is reset to none (N), since the prepare step has deleted the contents of the load and JCL datasets.

- 12. There is additional checking in the load (-I) step. The load step checks whether the schema currently in the database has been changed since the load files were prepared. If the schema has been changed, the load step is terminated to avoid loading entries prepared using an older schema that might not be valid for the newer schema.
- 13. During the check (-c) and prepare (-p) steps, the RDN of the new LDIF entry is checked to ensure that all the values specified in the RDN are also specified as attributes in the LDIF entry. Missing attributes are added to the entry before it is loaded into the directory. No messages are issued indicating this.
- 14. The Idif2tdbm program encrypts clear text userPassword attribute values for new entries loaded into the TDBM backend with the pwEncryption method specified in the configuration file. The Idif2tdbm program can load the LDIF format of an encrypted password unloaded by the tdbm2ldif program. The Idif2tdbm program cannot load the LDIF format unloaded by the tdbm2Idif program with the -t option. The -t option unloads encryption "tag visible" format passwords for use with non-z/OS LDAP servers. The **Idif2tdbm** program expects that textual data contained within the LDIF file is portable and of UTF-8 origin.
- 15. If you are maintaining multiple identical databases, you can use the prepare (-p) step of Idif2tdbm to prepare load data using one TDBM database and then invoke the load (-I) step of Idif2tdbm to load the data into a second TDBM database. When doing this, the contents of the second database must be identical to the first database at the time the load data was prepared, including the same schema entry in the database. After invoking the Idif2tdbm load (-I) step on the second database, you must also update the value of the NEXT_EID column in the DIR_MISC table in the second TDBM database. To do this, use the following SPUFI command to obtain the value of the NEXT_EID column in the first TDBM database after the **Idif2tdbm** prepare step is run:

SELECT NEXT EID FROM userid.DIR MISC;

Then use the following SPUFI command to update the NEXT_EID column in the second TDBM database to that value:

UPDATE userid.DIR MISC set NEXT EID = value;

- 16. Each loaded entry has a **creatorsname**, **modifiersname**, and **ibm-entryuuid** attribute. The values for these attributes can be specified in the LDIF input for the entry. If not:
 - The creatorsname and modifiersname attributes are assigned the value of the -b command line option if it is specified; otherwise, they are assigned the value of the adminDN option in the LDAP server configuration file.
 - The ibm-entryuuid attribute is assigned a unique value generated by the utility.

tdbm2ldif program

Purpose

This program is used to dump entries from a directory stored in a TDBM database into a file in LDAP Data Interchange Format (LDIF).

Format

```
tdbm21dif [-o outputFile] [-s subtreeDN | -n TdbmName] [-f confFile] [-d debugLevel] [-t]
```

Parameters

-o outputFile

Specify the output file to contain the directory entries in LDIF. All entries from the specified subtree are written in LDIF to the output file. If the file is not in the current directory, a fully-qualified name must be specified. If the **-o** option is not specified, then the output from the program is written to **stdout**.

-s subtreeDN

Identify the DN of the top entry of the subtree whose entries are to be converted to LDIF. This entry, plus all below it in the directory hierarchy, are converted and written to the output file. The -s option must be used to unload the schema entry. The -s option cannot be used when the -n option is specified.

-n TdbmName

Specify the name of the TDBM backend to unload. This is the name assigned to the backend on its database record in the configuration file. This can be used to indicate which TDBM backend to process when there are multiple TDBM backends in the configuration file. The -n option cannot be used when the **-s** option is specified.

Specify the name of the configuration file to use. This configuration file only needs to contain information for the TDBM backend which contains the entries to be converted to LDIF. Default is /etc/ldap/slapd.conf.

-d debugLevel

Specify the level of the debug messages to be created. The debug level is specified in the same fashion as the debug level for the LDAP server, as described on page 94. Table 16 on page 95 lists the specific debug levels. The default is no debug messages.

-t Specify that encrypted userPassword attribute values will be unloaded with their encryption tag in clear text, as follows:

```
userPassword: {tag}base64encoded and encryptedvalue
```

where tag is none, crypt, MD5, SHA, or DES: keylabel.

Examples

```
userPassword: {none}321p90fa0fdvn;a
userPassword: {crypt}3sdfaf[a
userPassword: {SHA}24309gf[jgt
userPassword: {DES:kgup.data.key}3ajewomv..=
```

In this format, the tag is visible, and only the userPassword value itself is encrypted and base64 encoded.

tdbm2ldif

Notes for using the -t parameter:

- 1. The format of data produced when -t is specified may be acceptable for other LDAP providers to load into their LDAP directory. And, if it is not directly loadable, this format is easily modified for loading by another provider into its LDAP directory. This format cannot be loaded back into an z/OS LDAP server.
- 2. The tag is enclosed by a left brace and a right brace. One colon is used between the userPassword keyword and the value, as opposed to two colons in the standard LDIF format of userPassword dumped by tdbm2ldif. This format cannot be read by the ldif2tdbm utility. It is intended for other LDAP providers and tools that may require the encryption tag visible.
- 3. Clear text passwords without a tag could still exist in the TDBM backend if the password was not modified or pwEncryption was not configured on the server. The values would be unloaded as standard binary attributes in base64 encoding. Following is an example:

userPassword:: kfa6903axs

4. The values returned by the crypt() algorithm are not portable to other X/Open-conformant systems. This means that user password values encoded by the crypt() algorithm and unloaded as tagged output using tdbm2ldif -t are not portable when loaded by another platform's load utility.

If -t is not specified, the output for the userPassword attribute is in the format:

userPassword::base64encodedValue

where base64encodedValue is a base64 encoded value of the binaryvalue, and binaryvalue={tag}encryptedPasswordValue.

All other command line inputs will result in a syntax error message, after which the proper syntax will be displayed. Also, specifying the same option multiple times with different values will result in a syntax error message.

Examples

Following is an example using the tdbm2ldif utility:

tdbm2ldif -o /tdbmdata/ldif.data -n tdbm1 -f /ldap/conf/slapd.conf

This tdbm2ldif utility invocation dumps all the data, except the schema entry, from the TDBM backend named tdbm1 in the configuration file to the file /tdbmdata/ldif.data, and uses the /ldap/conf/slapd.conf configuration file.

Usage

- 1. If the tdbm2ldif program is invoked with neither the -s nor the -n option and there is a single TDBM backend in the configuration file, all the directory entries in that TDBM databases are processed, except for the schema entry. If there are multiple TDBM backends in the configuration file, the program returns an error message.
 - 2. The tdbm2ldif program only dumps owner and ACL information for entries that have a specific owner or ACL. Any entry data with an inherited owner or ACL will not have owner or ACL information dumped.
 - 3. For the LDAP Version 3 protocol, there is a related set of Internet Drafts which discuss the introduction of a version mechanism for use in creating LDIF files. The tdbm2ldif utility always creates "tagged" LDIF files. The LDIF tag consists of a single line at the top of the LDIF file:

version: 1

All characters contained in the version: 1 LDIF file are portable characters represented in the local codepage. Strings containing nonportable characters (for instance, textual values containing multi-byte UTF-8 characters) must be base64 encoded.

- 4. To unload the schema entry from a TDBM directory, specify
 - -s cn=schema, suffixDN

- where suffixDN is the DN of a suffix in the TDBM directory. The schema entry is unloaded in LDIF modify format, which can be used in the -s option of the ldif2tdbm utility.
- 5. When editing the output file produced by tdbm2ldif, make sure to use an editor that does not delete blanks at the end of a line. If the output file contains a line that ends with blanks, using such an editor will result in deleting the blanks, thus changing the value of the attribute. This is especially serious when the line is continued, since the ending blanks will no longer be there to separate the last word in the line from the continuation on the next line. The maximum line length in a tdbm2ldif output file is 77; continued lines are always 77 characters long when the file is created by tdbm2ldif.
 - The **oedit** editor is an example of an editor that deletes blanks and should **not** be used.
- 6. The LDAP_DEBUG environment variable can also be used to set the debug level for tdbm2ldif. See page 94 for more information on specifying the debug level.

tdbm2ldif

Chapter 12. Running and using the LDAP entry UUID utility

The **Idapadduuids** utility adds **ibm-entryuuid** attributes to each entry in an LDAP TDBM directory. The **ibm-entryuuid** attribute holds a DCE standard Universally Unique Identifier (UUID). The UUID that is created is unique for all UUIDS on all computers.

Running the Idapadduuids utility from JCL

Sample JCL for running **Idapadduuids** from batch is provided with the LDAP server. The JCL includes an inline procedure, which will need to be modified by each installation to ensure that the **Idapadduuids** load module can be found. It may also be necessary to modify the JOB card for installation-specific requirements. "Sample JCL for Idapadduuids" on page 419 shows each sample JCL, which includes instructions for passing parameters into the utility programs. These jobs can be run by editing the JCL member of the *GLDHLQ*.SGLDSAMP dataset and entering the submit command. In addition, the **NLSPATH** variable must be set.

Running the Idapadduuids utility in the z/OS shell

In order to run the **Idapadduuids** utility in the shell, some environment variables need to be set properly. Sample shell scripts are shipped with the LDAP server for **Idapadduuids**. The sample shell scripts are located in **/usr/lpp/Idap/sbin**. These shell scripts need to be modified to fit your environment. The PATH, STEPLIB, and LANG variables should be set in the shell scripts. Ensure that **/usr/bin** is included in the **PATH** environment variable. Also, make sure **STEPLIB** is set to **GLDHLQ**.SGLDLNK if this PDS has not been added to LINKLIST, where **GLDHLQ** is replaced by the high level qualifier for the LDAP dataset. In addition, the **NLSPATH** variable must be set.

Idapadduuids utility

Purpose

The Idapadduuids utility adds ibm-entryuuids to each entry in a LDAP directory. This utility allows you to migrate all or portions of an existing directory to contain UUIDs. Directory applications may make use of the UUID to identify entries.

Format

Idapadduuids [options] [filter]

Parameters

options

The following table shows the options you can use for the Idapadduuids utility:

-?	Print this text.
-S method or	Specify the bind method to use. You can use either -m or -S to indicate the bind method
- m method	The default is SIMPLE . You can also specify GSSAPI to indicate a Kerberos Version 5 is requested, EXTERNAL to indicate that a certificate (SASL external) bind is requested, CRAM-MD5 to indicate that a SASL Challenge Response Authentication Mechanism bind is requested, or DIGEST-MD5 to indicate a SASL digest hash bind is requested.
	The GSSAPI method requires a protocol level of 3 and the user must have a valid Kerberos Ticket Granting Ticket in their credentials cache by using the Kerberos kinit command line utility.
	The EXTERNAL method requires a protocol level of 3. You must also specify -Z , -K , and -P to use certificate bind. If there is more than one certificate in the key database file, us -N to specify the certificate or the default certificate will be used.
	The CRAM-MD5 method requires protocol level 3. The -D or -U option must be specified
	The DIGEST-MD5 method requires protocol level 3. The -U option must be specified. The -D option can optionally be used to specify the authorization DN.
-h Idaphost	Specify the host on which the LDAP server is running. The default is the local host.
	When the target host is a z/OS LDAP server operating in multi-server mode with dynamic workload management enabled (see), the <i>Idaphost</i> value should be in the form <i>group_name.sysplex_domain_name</i> , where <i>group_name</i> is the name of the sysplexGroupName identified in the server configuration file and <i>sysplex_domain_name</i> is the name or alias of the sysplex domain in which the target server operates.
-p Idapport	Specify the TCP port where the LDAP server is listening. The default LDAP non-secure port is 389 and the default LDAP secure port is 636.
-d debuglevel	Set the LDAP debug level. The debug level you specify must be a decimal value. You ca also specify the levels additively if you want to specify two or more. For example, specify 32768 to turn on ERROR and specify 1 to turn on TRACE, or specify 32769 (32768+1) to turn on both ERROR and TRACE. See Table 16 on page 95 for the possible decimal values for <i>debuglevel</i> .
- D binddn	Use <i>binddn</i> to bind to the LDAP directory. The <i>binddn</i> parameter should be a string-represented DN. The default is a NULL string.
	If the -S or -m option is equal to DIGEST-MD5 or CRAM-MD5 , this option is the authorization DN which will be used for making access checks. This directive is optional when used in this manner.

-b searchbase	Use <i>searchbase</i> as the starting point for the search instead of the default. If -b is not specified, this utility examines the LDAP_BASEDN environment variable for a searchbase definition.
	If you are running in TSO, set the LDAP_BASEDN environment variable using LE runtime environment variable _CEE_ENVFILE . See <i>z/OS: C/C++ Programming Guide</i> for more information.
	If you are running in the z/OS shell, simply export the LDAP_BASEDN environment variable.
-Z	Use a secure connection to communicate with the LDAP server. Secure connections expect the communication to begin with the SSL/TLS handshake.
	The -K (<i>keyfile</i>) option or equivalent environment variable is required when the -Z option is specified. The -P (<i>keyfilepw</i>) option is required when the -Z option is specified and the key file specifies an HFS key database file. The -N (<i>keyfilelabel</i>) option must be specified if you wish to use a certificate that is different than the default specified in the key database.
-K keyfile	Specify the name of the System SSL key database file or RACF key ring. If this option is not specified, this utility looks for the presence of the SSL_KEYRING environment variable with an associated name.
	System SSL assumes that the name specifies a key database file. If the name is not a fully-qualified file name, then the current directory is assumed to contain the file. The key database file must be a file and cannot be an MVS dataset. If a corresponding file is not found then the name is assumed to specify a RACF key ring.
	See "SSL/TLS information for LDAP utilities" on page 145 for information on System SSL key databases and RACF key rings.
	This parameter is ignored if -Z is not specified.
-P keyfilepw	Specify either the key database file password or the file specification for a System SSL password stash file. When the stash file is used, it must be in the form file:// followed immediately (no blanks) by the HFS file specification (e.g. file:///etc/ldap/sslstashfile). The stash file must be a file and cannot be an MVS dataset.
	This parameter is ignored if -Z is not specified.
-N keyfiledn	Specify the certificate name in the key database file.
-U userName	Specify the <i>userName</i> for CRAM-MD5 or DIGEST-MD5 binds. The <i>userName</i> is a short name (uid) that will be used to perform bind authentication.
	This option is required if the -S or -m option is set to DIGEST-MD5.
-g realmName	Specify the <i>realmName</i> to use when doing a DIGEST-MD5 bind. This option is required when multiple realms are passed from an LDAP server to a client as part of a DIGEST-MD5 challenge; otherwise, it is optional.

filter

Specify an IETF RFC 1558 compliant LDAP search filter. UUID will be added only to those entries that match the filter.

Examples

Following are some **Idapadduuids** examples.

 $\bullet \quad \texttt{ldapadduuids -D "cn=ldap administrator" -w password -b "ou=Home Town,o=ibm_us,c=us" "cn=H*" } \\$ GLD2104I ldapadduuids added ibm-entryuuid to 1 entries.

This example adds UUIDs to all entries starting with "H" in the ou=Home Town subtree that do not already have an ibm-entryuuid attribute. In the sample data base this would be Henry Nguyen.

- ldapadduuids -D "cn=ldap administrator" -w password -b "ou=Home Town,o=ibm us,c=us" GLD2104I ldapadduuids added ibm-entryuuid to 51 entries.
 - This example adds UUIDs to all entries in the TDBM backend that did not already have one.

Usage

- 1. This utility only works with a TDBM backend.
- 2. With this utility you can break up the task of adding UUIDs to the database by specifying a subset of the database to add them to. This is useful when the directory database is very large and adding the UUIDs would take more time than your administrator's change window. It is also useful when you know that a portion of the directory tree will not need the entry UUIDs, thus saving space.
- 3. Note that an ibm-entryuuid attribute will not be added to entries that already have one.
- 4. New entries will automatically get an ibm-entryuuid attribute.

SSL/TLS information for LDAP utilities

The contents of a client's key database file is managed with the **gskkyman** utility. See *z/OS System* Secure Sockets Layer Programming Guide for information about the gskkyman utility. The gskkyman utility is used to define the set of trusted certification authorities (CAs) that are to be trusted by the client. By obtaining certificates from trusted CAs, storing them in the key database file, and marking them as trusted, you can establish a trust relationship with LDAP servers that use certificates issued by one of the CAs that are marked as trusted.

If the LDAP servers accessed by the client use server authentication, it is sufficient to define one or more trusted root certificates in the key database file. With server authentication, the client can be assured that the target LDAP server has been issued a certificate by one of the trusted CAs. In addition, all LDAP transactions that flow over the SSL/TLS connection with the server are encrypted, including the LDAP credentials that are supplied on the Idap_bind API.

For example, if the LDAP server is using a high-assurance VeriSign certificate, you should obtain a CA certificate from VeriSign, receive it into your key database file, and mark it as trusted. If the LDAP server is using a self-signed gskkyman server certificate, the administrator of the LDAP server can supply you with a copy of the server's certificate request file. Receive the certificate request file into your key database file and mark it as trusted.

Using the LDAP utilities without the -Z parameter and calling the secure port on an LDAP server (in other words, a non-secure call to a secure port) is not supported. Also, a secure call to a non-secure port is not supported.

SSL/TLS encrypts the keyring file. Either the password must be specified as part of the -P parameter or file specification of a stash file that was created using the gskkyman utility must be specified in the form file://followed immediately (no blanks in between) by the file specification of the stash file.

Using RACF key rings

Alternately, LDAP supports the use of a RACF key ring. See the certificate/key management section in z/OS: System Secure Sockets Layer Programming for instructions on how to migrate a key database to RACF and how to use the **RACDCERT** command to protect the certificate and key ring.

The user ID under which the LDAP client runs must be authorized by RACF to use RACF key rings. To authorize the LDAP client, you can use the RACF commands in the following example (where userid is the user ID running the LDAP client utility):

RDEFINE FACILITY IRR.DIGTCERT.LIST UACC(NONE) RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(NONE) PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(userid) ACCESS(CONTROL) PERMIT IRR.DIGTCERT.LIST CLASS(FACILITY) ID(userid) ACCESS(CONTROL)

Remember to refresh RACF after doing the authorizations.

| SETROPTS RACLIST(FACILITY) REFRESH

Once the RACF key ring is set up and authorized, specify the RACF key ring name for the -K keyfile option and do not specify the **-P** *keyfilepw* option.

Chapter 13. Running and using the password encryption utility

The **db2pwden** utility is an administration utility used to migrate clear text user passwords in the TDBM backend to encrypted passwords. The **db2pwden** utility can be run the z/OS shell or from TSO.

Running the db2pwden utility from JCL

Sample JCL for running **db2pwden** from batch is provided with the LDAP server. The JCL includes an inline procedure, which will need to be modified by each installation to ensure that the **db2pwden** load module can be found. It may also be necessary to modify the JOB card for installation-specific requirements. Appendix E, "Sample JCL" on page 413 shows each sample JCL, which includes instructions for passing parameters into the utility programs. These jobs can be run by editing the JCL member of the *GLDHLQ*.SGLDSAMP dataset and entering the submit command. In addition, the **NLSPATH** variable must be set.

Running the db2pwden utility in the z/OS shell

In order to run the **db2pwden** utility in the shell, some environment variables need to be set properly. Sample shell scripts are shipped with the LDAP server for **db2pwden**. The sample shell scripts are located in **/usr/lpp/ldap/sbin**. These shell scripts need to be modified to fit your environment. The PATH, STEPLIB, and LANG variables should be set in the shell scripts. Ensure that **/usr/bin** is included in the **PATH** environment variable. Also, make sure **STEPLIB** is set to **GLDHLQ**.SGLDLNK if this PDS has not been added to LINKLIST, where **GLDHLQ** is replaced by the high level qualifier for the LDAP dataset. In addition, the **NLSPATH** variable must be set.

db2pwden utility

Purpose

The db2pwden utility is provided to encrypt all clear text user passwords in an already loaded TDBM backend. The utility runs as a client operation while the server is active, and causes the server to encrypt all the userPassword attribute values that are in clear text with the pwEncryption method configured on the LDAP server. The utility must be run by the LDAP administrator.

Format

db2pwden [options]

Parameters

options

The following table shows the *options* you can use for the **db2pwden** utility:

Table 21. db2pwden options

Option	Description
-?	Print this text.
-m mechanism	Specify the bind method to use. You can use either -m or -S to indicate the bind method.
or -S mechanism	The default is SIMPLE . You can also specify GSSAPI to indicate a Kerberos Version 5 is requested, EXTERNAL to indicate that a certificate (SASL external) bind is requested, CRAM-MD5 to indicate that a SASL Challenge Response Authentication Mechanism bind is requested, or DIGEST-MD5 to indicate a SASL digest hash bind is requested.
	The GSSAPI method requires a protocol level of 3 and the user must have a valid Kerberos Ticket Granting Ticket in their credentials cache by using the Kerberos kinit command line utility.
	The EXTERNAL method requires a protocol level of 3. You must also specify -Z , -K , and -P to use certificate bind. If there is more than one certificate in the key database file, use -N to specify the certificate or the default certificate will be used.
	The CRAM-MD5 method requires protocol level 3. The -D or -U option must be specified.
	The DIGEST-MD5 method requires protocol level 3. The -U option must be specified. The -D option can optionally be used to specify the authorization DN.
-h Idaphost	Specify the host on which the LDAP server is running. The default is the local host.
	When the target host is a z/OS LDAP server operating in multi-server mode with dynamic workload management enabled (see "Determining operational mode" on page 75 for additional information about LDAP server operating modes), the <i>Idaphost</i> value should be in the form <i>group_name.sysplex_domain_name</i> , where <i>group_name</i> is the name of the sysplexGroupName identified in the server configuration file and <i>sysplex_domain_name</i> is the name or alias of the sysplex domain in which the target server operates.
-p Idapport	Specify the TCP port where the LDAP server is listening. The default LDAP non-secure port is 389 and the default LDAP secure port is 636.
-d debuglevel	Specify the level of debug messages to be created. The debug level is specified in the same fashion as the debug level for the LDAP server, as described on page 94. Table 16 on page 95 lists the specific debug levels. The default is no debug messages.

Table 21. db2pwden options (continued)

Option	Description
-D binddn	Use <i>binddn</i> to bind to the LDAP directory. The <i>binddn</i> parameter should be a string-represented DN. The default is a NULL string.
	If the -S or -m option is equal to DIGEST-MD5 or CRAM-MD5 , this option is the authorization DN which will be used for making access checks. This directive is optional when used in this manner.
-w bindpasswd	Use bindpasswd as the password for simple authentication. The default is a NULL string.
-b base	Use base as the starting point for the update instead of the default. If -b is not specified, thi utility examines the LDAP_BASEDN environment variable for a base definition.
	If you are running in TSO, set the LDAP_BASEDN environment variable using LE runtime environment variable _CEE_ENVFILE . See <i>z/OS: C/C++ Programming Guide</i> for more information.
	If you are running in the z/OS shell, export the LDAP_BASEDN environment variable.
-Z	Use a secure connection to communicate with the LDAP server. Secure connections expect the communication to begin with the SSL/TLS handshake.
	The -K (<i>keyfile</i>) option or equivalent environment variable is required when the -Z option is specified. The -P (<i>keyfilepw</i>) option is required when the -Z option is specified and the key file specifies an HFS key database file. The -N (<i>keyfilelabel</i>) option must be specified if you wish to use a certificate that is different than the default specified in the key database.
-K keyfile	Specify the name of the System SSL key database file or RACF key ring. If this option is no specified, this utility looks for the presence of the SSL_KEYRING environment variable with an associated name.
	System SSL assumes that the name specifies a key database file. If the name is not a fully-qualified file name, then the current directory is assumed to contain the file. The key database file must be a file and cannot be an MVS dataset. If a corresponding file is not found then the name is assumed to specify a RACF key ring.
	See "SSL/TLS information for LDAP utilities" on page 145 for information on System SSL key databases and RACF key rings.
	This parameter is ignored if -Z is not specified.
-P keyfilepw	Specify either the key database file password or the file specification for a System SSL password stash file. When the stash file is used, it must be in the form file:// followed immediately (no blanks) by the HFS file specification (e.g. file:///etc/ldap/sslstashfile). The stash file must be a file and cannot be an MVS dataset.
	This parameter is ignored if -Z is not specified.
-N keyfiledn	Specify the label associated with the key in the key database file.
-U username	Specify the <i>userName</i> for CRAM-MD5 or DIGEST-MD5 binds. The <i>userName</i> is a short name (uid) that will be used to perform bind authentication.
	This option is required if the -S or -m option is set to DIGEST-MD5.
-g realmname	Specify the <i>realmName</i> to use when doing a DIGEST-MD5 bind. This option is required when multiple realms are passed from an LDAP server to a client as part of a DIGEST-MD challenge; otherwise, it is optional.

Examples

Following are some db2pwden examples:

• The following command: db2pwden -D "cn=admin" -w secret

db2pwden

- encrypts all clear text user passwords in the TDBM backend at the LDAP server on the local host. The encryption method used is the **pwEncryption** method configured on the LDAP server.
- The following command:

```
db2pwden -h ushost -p 391 -D "cn=admin" -w secret -b "o=university, c=US"
```

encrypts all clear text user passwords starting at the base "o=university,c=US" in the TDBM backend on host ushost at port 391. The encryption method used is the **pwEncryption** method configured on the LDAP server.

SSL/TLS notes

The contents of a client's key database file is managed with the **gskkyman** utility. See *z/OS: System* Secure Sockets Layer Programming for information about the gskkyman utility. The gskkyman utility is used to define the set of trusted certification authorities (CAs) that are to be trusted by the client. By obtaining certificates from trusted CAs, storing them in the key database file, and marking them as trusted, you can establish a trust relationship with LDAP servers that use certificates issued by one of the CAs that are marked as trusted.

If the LDAP servers accessed by the client use server authentication, it is sufficient to define one or more trusted root certificates in the key database file. With server authentication, the client can be assured that the target LDAP server has been issued a certificate by one of the trusted CAs. In addition, all LDAP transactions that flow over the SSL/TLS connection with the server are encrypted, including the LDAP credentials that are supplied on the Idap_bind API.

For example, if the LDAP server is using a high-assurance VeriSign certificate, you should obtain a CA certificate from VeriSign, receive it into your key database file, and mark it as trusted. If the LDAP server is using a self-signed gskkyman server certificate, the administrator of the LDAP server can supply you with a copy of the server's certificate request file. Receive the certificate request file into your key database file and mark it as trusted. If the LDAP server accessed by the client is configured to server and client authentication, then the client can transmit its certificate for authentication by the server by marking it as the default in its key database file.

Using this utility without the -Z parameter and calling the secure port on an LDAP server (a non-secure call to a secure port) is not supported.

System SSL encrypts the key ring file. Either the password must be specified as part of the -P parameter or file specification of a stash file that was created using the **gskkyman** utility must be specified in the form "file://" followed immediately (no blanks in between) by the file specification of the stash file.

Diagnostics

Exit status is 0 if no errors occur. Errors result in a non-zero exit status and a diagnostic message being written to standard error.

Chapter 14. Internationalization support

This chapter discusses translated messages and UTF-8 support.

Translated messages

The **LANG** and **NLSPATH** environment variables are set for the LDAP Server and the LDAP programs in the server's **envvars** file:

/etc/ldap/slapd.envvars

There are no default values for these variables. Messages are also available in Japanese. The **LANG** variable should be set to **LANG=Ja_JP** or **Ja_JP.IBM-939**. These variables should also be set either in the environment variable file of the user or by exporting the variables in the shell for the user ID that will run the LDAP utilities.

Following is part of a sample **slapd.envvars** file.

LANG=En_US.IBM-1047 NLSPATH=/usr/lib/nls/msg/%L/%N:/usr/lib/nls/msg/En_US.IBM-1047/%N

- There are symbolic links to the English language message catalogs in the following locations:
- /usr/lib/nls/msg/C
- /usr/lib/nls/msg/En_US
- /usr/lib/nls/msg/En_US.IBM-1047
- It is possible to run with either LANG=En_US, LANG=C, or LANG=En_US.IBM-1047 and access the English language LDAP message catalogs.
 - There are also symbolic links to the following Japanese language message catalogs:
- /usr/lib/nls/msg/Ja_JP
- /usr/lib/nls/msg/Js JP.IBM-939

UTF-8 support

UTF stands for "UCS (Unicode) Transformation Format". The UTF-8 encoding can be used to represent any Unicode character. Depending on a Unicode character's numeric value, the corresponding UTF-8 character is a 1, 2, or 3 byte sequence. Table 22 shows the mapping between Unicode and UTF-8. Refer to IETF RFC 2279 UTF-8, a transformation format of ISO 10646 for more information on UTF-8.

Table 22. Mapping between Unicode and UTF-8

Unicode range (hexadecimal)	UTF-8 octet sequence (binary)
0000-007F	0xxxxxxx
0080-07FF	110xxxxx 10xxxxxx
0800-FFFF	1110xxxx 10xxxxxx 10xxxxxx

The LDAP Version 3 protocol specifies that all data exchanged between LDAP clients and servers be UTF-8. The z/OS LDAP server supports UTF-8 data exchange as part of its Version 3 protocol support.

Note: For UTF-8 data stored in a z/OS LDAP server's TDBM backends, collation for single-byte UTF-8 characters is relative to the server's locale. For multi-byte UTF-8 characters, collation is relative to the numeric value of the equivalent Unicode character.

Part 2. Use

Chapter 15. Data model

The LDAP data model is closely aligned with the X.500 data model. In this model, a directory service provides a hierarchically organized set of *entries*. Each of these entries is represented by an *object class*. The object class of the entry determines the set of *attributes* which are required to be present in the entry as well as the set of attributes that can optionally appear in the entry. An attribute is represented by an *attribute type* and one or more *attribute values*. In addition to the attribute type and values, each attribute has an associated *syntax* which describes the format of the attribute values. Examples of attribute syntaxes for LDAP include directory string and **binary**.

To summarize, the directory is made up of entries. Each entry contains a set of attributes. These attributes can be single or multi-valued (have one or more values associated with them). The object class of an entry determines the set of attributes that must and the set of attributes that may exist in the entry. Refer to for more information on the X.500 Directory Information Model.

Every entry in the directory has a *distinguished name (DN)*. The DN is the name that uniquely identifies an entry in the directory. A DN is made up of attribute=value pairs, separated by commas. For example:

```
cn=Ben Gray,ou=editing,o=New York Times,c=US
cn=Lucille White,ou=editing,o=New York Times,c=US
cn=Tom Brown,ou=reporting,o=New York Times,c=US
```

- Any non-binary attributes defined in the directory schema may be used to make up a DN.
- Note: The z/OS LDAP server does not allow a DN that contains a userpassword attribute.

The order of the component attribute=value pairs is important. The DN contains one component for each level of the directory hierarchy. LDAP DNs begin with the most specific attribute (usually some sort of name), and continue with progressively broader attributes, often ending with a country attribute.

Relative distinguished names

Each component of a DN is referred to as a *relative distinguished name (RDN)*. It identifies an entry distinctly from any other entries which have the same parent. In the examples above, the RDN cn=Ben Gray separates the first entry from the second entry, (with RDN cn=Lucille White). The attribute=value pair or pairs making up the RDN for an entry must also be present as an attribute=value pair or pairs in the entry. This is not true of the other components of the DN. When using the TDBM backend, TDBM adds the attribute=value pairs in the RDN to the entry if they are not already present.

RDNs can contain multiple attribute=value pairs. So-called multivalued RDNs use two or more attribute:value pairs from the directory entry to define the name of the entry relative to its parent. An example where this would be useful would be where a directory hierarchy of users was being defined for a large university. This hierarchy would be segmented by campus. A problem is encountered, however, when it is discovered that there is more than one John Smith at the downtown campus. The RDN cannot simply be the name of the user. What can be done, however, is add a unique value to the RDN, thus ensuring its uniqueness across the campus. Typically universities hand out serial numbers to their students. Coupling the student number with the person's name is one method of solving the problem of having a unique RDN under a parent in the directory hierarchy. The entry's RDN might look something like:

cn=John Smith+studentNumber=123456.

The plus sign (+) is used to delimit separate attribute=value pairs within an RDN. The entry's DN might look like:

cn=John Smith+studentNumber=123456, ou=downtown, o=Big University, c=US

	Distinguished name syntax
	The Distinguished Name (DN) syntax supported by this server is based on IETF RFC 1779 A String Representation of Distinguished Names and IETF RFC 2253 LDAP (v3): UTF-8 String Representation of Distinguished Names. A semicolon (;) character may be used to separate RDNs in a distinguished name, although the comma (,) character is the typical notation.
	White space (blank) characters may be present on either side of the comma or semicolon. The white space characters are ignored, and the semicolon replaced with a comma.
	In addition, space characters may be present between an attribute=value pair and a plus sign (+), between an attribute type and an equal sign (=), and between an equal sign (=) and an attribute value. These space characters are ignored when parsing.
	A value may be surrounded by quotation marks, which are not part of the value. Inside the quoted value, the following characters can occur without any escaping: • A space or pound sign (#) character occurring at the beginning of the string • One of the characters - apostrophe (') - equal sign (=) - plus sign (+) - backslash (\) - less than sign (<) - greater than sign (>) - semicolon (;)
	Alternatively, a single character to be escaped may be prefixed by a backslash (\). This method may be used to escape any of the characters listed above, plus the quotation mark.
 	This notation is designed to be convenient for common forms of name. This section gives a few examples of distinguished names written using this notation. First is a name containing three components: OU=Sales+CN=J. Smith,0=Widget Inc.,C=US
	This example shows a method of escaping a comma in an organization name: CN=R. Smith,0=Big Company Inc.,C=US
	Domain component naming Domain component naming as specified by RFC 2247 is also supported in the LDAP server. For example, the domain name <code>ibm.com</code> could be specified as an entry in the LDAP server with the following distiguished name: dc=ibm,dc=com
	RACF-style distinguished names If you are using SDBM (the RACF database backend of the LDAP server), the format of the DNs is restricted in order to match the schema of the underlying RACF data. A RACF-style DN for a user or group contains two required attributes plus a suffix:

racfid Specifies the user ID or group ID.

profiletype

Specifies user or group.

suffix Specifies the SDBM suffix.

A RACF-style DN for a user's connection to a group also contains two required attributes plus a suffix:

```
racfuserid+racfgroupid
          Specifies the user and the group.
  profiletype
          Specifies connect.
  suffix Specifies the SDBM suffix.
  The suffix for SDBM may contain additional attributes. For example, if the suffix has been specified as:
  suffix cn=myRACF,c=US
  in the LDAP configuration file, any RACF-style DN would end with:
cn=myRACF,c=US
| Following is an example of the DN format and a sample DN for a user:
| racfid=userid,profiletype=user,suffix
| racfid=ID1,profiletype=user,cn=myRACF,c=US
| Following is an example of the DN format and a sample DN for a connection:
| racfuserid=userid+racfgroupid=groupid,profiletype=connect,suffix
  racfuserid=ID1+racfgroupid=GRP1,profiletype=connect,cn=myRACF,c=US
```

Chapter 16. LDAP directory schema for TDBM

The LDAP Version 3 (V3) protocol, as defined in IETF RFC 2252 Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions and IETF RFC 2256 A Summary of the X.500(96) User Schema for use with LDAPv3, describes schema publication and update. Schema publication provides the ability to query the active directory schema through the use of the LDAP search function. Schema update is the ability to change the schema while the directory server is running.

Note: The z/OS LDAP server implementation of both schema publication and update is provided for the TDBM backend only. Schema publication only is available for the SDBM backend.

Setting up the schema for TDBM - new users

The LDAP server is shipped with two predefined schema files representing schema definitions which the user may want to load as the LDAP schema TDBM. For TDBM, the schema is stored as an entry in the database and modify operations may be performed on this entry. These files are **schema.user.Idif** and **schema.IBM.Idif**. The **schema.IBM.Idif** schema definitions require that the definitions contained in **schema.user.Idif** are loaded prior to loading **schema.IBM.Idif**. Determine which schema files would be used to represent the data stored in the TDBM database. Copy the files from the **/usr/lpp/Idap/etc** directory to a working directory, for example, the **/home/myuser** directory. For each file, find the line dn: cn=schema, <suffix>

and replace *suffix*> with one of the suffixes defined for TDBM in the LDAP server configuration file. In the example below, the suffix defined in the configuration file for TDBM is o=YourCompany,c=US. Then run the **Idapmodify** command from the z/OS shell specifying the host, port, bind DN, password, and schema file for each schema file to be loaded. This will load the schema into the directory.

For example:

- 1. Copy the schema file(s) to a working directory:
 - cp /usr/lpp/ldap/etc/schema.user.ldif /home/myuser/schema.user.ldif
 - cp /usr/lpp/ldap/etc/schema.IBM.ldif /home/myuser/schema.IBM.ldif
- 2. Edit the schema file(s) and replace

```
dn: cn=schema, <suffix>
```

with

dn:cn=schema, o=Your Company, c=US

| 3. Run the **Idapmodify** command(s):

```
ldapmodify -h ldaphost -p ldapport -D adminDN -w passwd -f /home/myuser/schema.user.ldif ldapmodify -h ldaphost -p ldapport -D adminDN -w passwd -f /home/myuser/schema.IBM.ldif
```

See *z/OS: Security Server LDAP Client Programming* for more information about **Idapmodify**.

Upgrading schema for TDBM

The schema files that are shipped with the z/OS LDAP server are based on industry and product defined schemas. As such, they should not be modified since existing products and applications use the schema elements as defined.

Prior to z/OS V1R4, individual schema files related to specific uses were shipped with the LDAP server to be used for TDBM. With APAR OW53447, users of the TDBM backend were requested to move from using individual schema files to only use the **schema.user.ldif** or **schema.user.ldif** plus **schema.lBM.ldif** files for TDBM. Customers using TDBM for the first time must only use **schema.user.ldif** or **schema.user.ldif** plus **schema.lBM.ldif**.

The individual schema files previously shipped with the LDAP server have been shipped with z/OS V1R4 LDAP solely for the purpose of upgrading the schema to be equivalent to the schema.user.ldif and schema.IBM.Idif schema files. Only use the individual schema files if you are directed to use them.

Determine which of the following cases applies to your LDAP installation:

- 1. If the TDBM backend has never been configured and started, refer to "Setting up the schema for TDBM - new users" on page 159 for instructions on loading the schema for TDBM.
- 2. If you are migrating fron an RDBM backend to a TDBM backend, refer to "Migrating the schema from RDBM to TDBM" on page 176.
- 3. If you are currently running OS/390 V2R10, z/OS V1R1, z/OS V1R2, or z/OS V1R3 and APAR OW53447 has been applied and the schema update instructions have been completed, refer to "Updates to the schema".
- 4. If you are currently running OS/390 V2R10, z/OS V1R1, z/OS V1R2, or z/OS V1R3 and have not applied the schema modifications from APAR OW53447, refer to Appendix I, "TDBM schema migration" on page 439.

Updates to the schema

Occasionally, schema updates are required during the life of an LDAP release. These updates are shipped in a file named schema-updates.ldif located in the /usr/lpp/ldap/etc directory. Specific instructions for applying the schema updates are included as comments in the file. When moving to a new LDAP release, all updates that have been provided through the schema-updates.ldif file must be applied to your schema for the schema modules which you have loaded into your TDBM schema. Future schema modules will depend on those updates being applied to your schema.

When the schema-updates.ldif file is shipped through the service stream, the administrator should review the comments in the file to determine if any of the changes apply to their TDBM schema. If the changes apply, the administrator should:

- 1. Copy the file to a local directory
- 2. Update the cn=schema, <suffix> line in the file for the modifications they want to apply
- 3. Remove or comment out the other modifications in the file
- 4. Apply the changes to their TDBM schema through the **Idapmodify** utility

Schema introduction

Entries in the directory are made up of attributes which consist of an attribute type and one or more attribute values. These are referred to as attribute=value pairs. Every entry contains one or more objectClass attributes that identify what type of information the entry contains. The object classes associated with the entry determine the set of attributes which must or may be present in the entry.

The schema is represented and stored as another entry in the directory. Following is a portion of the schema entry.

```
cn=SCHEMA,o=Your Company,c=US
subtreespecification=NULL
objectclass=TOP
objectclass=SUBSCHEMA
objectclass=SUBENTRY
objectclass=IBMSUBSCHEMA
attributetypes= ( 2.5.4.3 NAME ( 'cn' 'commonName' ) SUP name )
ibmattributetypes = ( 2.5.4.3 ACCESS-CLASS normal )
objectclasses = ( 2.5.6.0 NAME 'top' ABSTRACT MUST objectclass )
ldapsyntaxes = ( 1.3.6.1.4.1.1466.115.121.1.15 DESC 'directory string' )
matchingrules = ( 2.5.13.5 NAME 'caseExactMatch' SYNTAX '1.3.6.1.4.1.1466.115.121.1.15 )
```

Figure 16. Sample Schema Entry

The objectClass values specified for the schema entry are top, subEntry, subSchema, and ibmSubschema. This set of object classes result in the objectClass, cn, and subtreeSpecification attributes being required for a schema entry and the attributeTypes, objectClasses, IdapSyntaxes, matchingRules, and ibmAttributeTypes attributes being allowed in a schema entry.

Note: The ditContentRules, ditStructureRules, nameforms, and matchingRuleUse attributes are allowed in a schema entry, but usage of these directives is not implemented by the z/OS LDAP server.

Every entry in the directory including the schema entry contains the subschemaSubentry attribute. The value shown for this attribute is the DN of the schema entry which contains the definitions for the directory entry. For entries stored in TDBM, the **subschemaSubentry** value will be cn=schema, <suffix> where <suffix> is the TDBM suffix under which this directory entry is stored. For example, assuming the suffix is 'o=Acme Company, c=UK', for the directory entry

```
cn= Mary Smith, o=Acme Company, c=UK
objectClass= person
cn= Mary Smith
sn= Smith
```

the **subschemaSubentry** value would be

```
subschemaSubentry= cn=schema, o=Acme Company, c=UK
```

Attribute types, object classes, LDAP syntaxes, and matching rules have assigned unique numeric object identifiers. These numeric object identifiers are in dotted decimal format, for example, 2.5.6.6. Attribute types, object classes, and matching rules are also identified by a textual name, for example, person or names. The numeric object identifier and the textual names may be used interchangeably when an attribute type or object class definition specifies an object identifier. Most schema definitions use the textual name as the object identifier for these definitions.

The attributes that comprise a directory schema include attribute types, IBM attribute types, object classes, LDAP syntaxes, and matching rules. There is a fixed set of LDAP syntaxes and matching rules supported by the z/OS LDAP server. These are listed in Table 26 on page 168, Table 27 on page 168, and Table 28 on page 169. Each of the schema attributes are described below.

· Attribute types

Attribute types define the characteristics of the data values stored in the directory. Each attribute type defined in a schema must contain a unique numeric object identifier and optionally contain a textual name, zero or more alias names, and a description of the attribute type. The characteristics defined for each attribute type include the syntax, length, and matching rules.

The **SYNTAX** defines the format of the data stored for the attribute type. The server checks the attribute values that are to be added to the directory by comparing the values against the set of allowed characters based on the syntax. For example, if the syntax of an attribute type is Boolean (where the acceptable values are TRUE or FALSE) and the attribute value specified is ves, the update will fail. The syntaxes supported by the z/OS LDAP server are shown in Table 26 on page 168 and Table 27 on page 168.

Matching rules may be specified for **EQUALITY**, **ORDERING**, and **SUBSTR** (substring matching). The matching rule determines how comparisons between values will be done. The EQUALITY matching rule determines if two values are equal. Examples of **EQUALITY** matching rules are **caseIgnoreMatch**, caseExactMatch, and telephoneNumberMatch. The ORDERING matching rule determines how two values are ordered (greaterThanOrEqual, lessThanOrEqual). Examples of ORDERING matching rules are caselgnoreOrderingMatch and generalizedTimeOrderingMatch. The SUBSTR matching rule determines if the presented value is a substring of an attribute value from the directory. Examples of SUBSTR matching rules are caselgnoreSubstringsMatch and telephoneNumberSubstringsMatch.

If EQUALITY, ORDERING, or SUBSTR matching rules are not specified in the definition of an attribute type or through the inheritance hierarchy, the z/OS LDAP server will perform evaluations to the best of its ability, but the results may not be as expected. The z/OS LDAP server uses the matching rules shown in the following table based on attribute type syntax to evaluate EQUALITY if the EQUALITY matching rule is not specified.

Table 23. Syntax and default EQUALITY matching rule

Syntax	Equality
Attribute Type Description	objectIdentifierFirstComponentMatch
Binary	-
Boolean	caselgnoreMatch
Directory String	caselgnoreMatch
DIT Content Rule Description	objectIdentifierFirstComponentMatch
DIT Structure Rule Description	integerFirstComponentMatch
DN	distinguishedNameMatch
Generalized Time	generalizedTimeMatch
IA5 String	caselgnoreIA5Match
IBM Attribute Type Description	objectIdentifierFirstComponentMatch
IBM Entry UUID	IBM-EntryUUIDMatch
Integer	integerMatch
LDAP Syntax Description	objectIdentifierFirstComponentMatch
Matching Rule Description	objectIdentifierFirstComponentMatch
Matching Rule Use Description	objectIdentifierFirstComponentMatch
Name Form Description	objectIdentifierFirstComponentMatch
Object Class Description	objectIdentifierFirstComponentMatch
Object Identifier	caselgnoreMatch
Octet String	-
Substring Assertion	-
Telephone Number	telephoneNumberMatch
UTC Time	utcTimeMatch

The z/OS LDAP server also verifies that the matching rules specified for **EQUALITY**, **ORDERING**, and SUBSTR are consistent with the specified SYNTAX. Table 24 shows acceptable values EQUALITY, ORDERING, and SUBSTR.

Table 24. Syntax and acceptable matching rules (EQUALITY, ORDERING, and SUBSTR)

Syntax	EQUALITY	ORDERING	SUBSTR
Attribute Type Description	objectIdentifierFirstComponentMatch	-	-
Binary	-	-	-
Boolean	booleanMatch caseIgnoreMatch caseIgnoreIA5Match caseExactMatch caseExactIA5Match	caseIgnoreOrderingMatch caseExactOrderingMatch	caseIgnoreSubstringsMatch caseExactSubstringsMatch
Directory String	caseIgnoreMatch caseExactMatch	caseIgnoreOrderingMatch caseExactOrderingMatch	caseIgnoreSubstringsMatch caseExactSubstringsMatch
DIT Content Rule Description	objectIdentifierFirstComponentMatch	-	-
DIT Structure Rule Description	integerFirstComponentMatch	-	-
DN	distinguishedNameMatch	distinguishedNameOrderingMat	teh
Generalized Time	generalizedTimeMatch	generalizedTimeOrdering Match	-
IA5	caselgnoreMatch caselgnoreIA5Match caseExactMatch caseExactIA5Match	caseIgnoreOrderingMatch caseExactOrderingMatch	caseIgnoreSubstringsMatch caseExactSubstringsMatch
IBM Attribute Type Description	objectIdentifierFirstComponentMatch	-	-
IBM Entry UUID	IBM-EntryUUIDMatch	-	-
Integer	integerMatch	-	-
LDAP Syntax Description	objectIdentifierFirstComponentMatch	-	-
Matching Rule Description	objectIdentifierFirstComponentMatch	-	-
Matching Rule Use Description	objectIdentifierFirstComponentMatch	-	-
Name Form Description	object Identifier First Component Match	-	-
Object Class Description	objectIdentifierFirstComponentMatch	-	-
Object Identifier	objectIdentifierMatch	-	-
Octet String	octetStringMatch	-	-
Substring Assertion	-	-	-
Telephone Number	telephoneNumberMatch	-	telephoneNumberSubstrings Match

The syntax or matching rule values may be inherited by specifying a superior attribute type. This is done by specifying the keyword **SUP**, followed by the object identifier of the superior attribute type. This is known as an attribute type hierarchy and referred to as inheritance. A superior hierarchy may be created with multiple levels of inheritance. In the following partial example, ePersonName and personName would inherit their SYNTAX from name.

```
ePersonName SUP personName
personName SUP name
name SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
```

When the SYNTAX, EQUALITY, ORDERING, or SUBSTR values are not specified for an attribute type, the attribute type hierarchy will be used to determine these values. The SYNTAX must be specified on the attribute type or through inheritance.

The number of values that may be stored in each entry for an attribute type is limited to one value if the keyword **SINGLE-VALUE** is specified. Otherwise, any number of attribute values may exist in the entry.

The **OBSOLETE** keyword indicates that the attribute type cannot be used to add data to existing entries or to store data in new entries. Modifications to entries which contain data values of an attribute type which has been made obsolete will fail unless all data values for all obsolete attribute types are removed during the modification. Searches specifying the obsolete attribute type will return the entries containing the attribute type. If an obsolete attribute type is referred to in a superior hierarchy, the inherited values will continue to be resolved.

Example 1:

```
attributeTypes: ( 1.2.3.4 NAME 'obsattr1' SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 OBSOLETE )
attributeTypes: ( 5.6.7.8 NAME 'validattr1' SUP obsattr1 )
would be the same as
attributeTypes: ( 5.6.7.8 NAME 'validattr' SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )
Example 2:
attributeTypes: ( 10.20.30.40 NAME 'obsattr2' SUP obsattr3 )
attributeTypes: (50.60.70.80 NAME 'obsattr3'
EQUALITY caseIgnoreMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 )
attributeTypes: ( 90.100.110.120 NAME 'validattr2' SUP obsattr2 )
would be the same as
attributeTypes: ( 90.100.110.120 NAME 'validattr2'
EQUALITY caseIgnoreMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 )
```

The USAGE keyword's valid values are userApplications or one of three operational values (directoryOperation, distributedOperation, or dSAOperation). Only LDAP servers may define the attribute types which have an operational usage value. User schemas may only define attribute types with USAGE userApplications specified. The default is userApplications, therefore USAGE does not need to be specified in user schemas.

The z/OS LDAP server restricts users from modifying data values specified for an attribute type when NO-USER-MODIFICATION is specified on the definition of the attribute type. NO-USER-**MODIFICATION** may only be specified for operational attribute types.

Note: The LDAP V3 protocol also defines a COLLECTIVE key word for attribute types. The LDAP server does not support this key word. All attribute types are assumed to be not COLLECTIVE.

IBM attribute types

Additional information required by IBM LDAP servers for each attribute type defined in the schema is specified using the IBMAttributeTypes schema attribute. The IBMAttributeTypes schema attribute is an extension of the attributeTypes schema attribute. If the attributeTypes value is not defined, then

the corresponding IBMAttributeTypes value cannot be defined. For the z/OS LDAP server, the additional information defined using this attribute is the ACCESS-CLASS of the associated attribute type.

The ACCESS-CLASS specifies the level of access users have to data values of this attribute type. The levels that may be specified for user-defined attribute types are normal, sensitive, and critical. The system and restricted keywords are reserved for LDAP server use and are specified for some of the attribute types controlled by the server. See "Attribute access classes" on page 251 for definition of access classes.

Note: Other LDAP servers from IBM use the DBNAME and LENGTH characteristics to specify additional information for their implementations. These may be specified in the schema but are not used by the z/OS LDAP server.

· Object classes

Object classes define the characteristics of individual directory entries. The object classes listed in a directory entry determine the set of required and optional attributes for the entry. Each object class defined in a schema must contain a unique numeric object identifier and optionally contain a textual name, zero or more alias names, a description of the object class, and lists of required (MUST) or optional (MAY) attribute types.

Required and optional attribute types for an object class may be inherited by specifying one or more superior object classes in an object class definition. This is done by specifying the keyword SUP followed by the object identifiers of the superior object classes. This is known as an object class hierarchy and referred to as multiple inheritance. A superior hierarchy may be created with multiple levels of inheritance.

Each object class is defined as one of three types: STRUCTURAL, ABSTRACT, or AUXILIARY. The type is specified when the object class is defined. If the type is not specified, it defaults to STRUCTURAL.

The structural object class defines the characteristics of a directory entry. Each entry must specify exactly one base structural object class. A base structural object class is defined as the most subordinate object class in an object class hierarchy. The structural object class of an entry can not be changed. Once an entry is defined in the directory, it must be deleted and recreated to change the structural object class.

Abstract and auxiliary object classes are used to provide common characteristics to entries with different structural object classes. Abstract object classes are used to derive additional object classes. Abstract object classes must be referred to in a structural or auxiliary superior hierarchy. Auxiliary object classes are used to extend the set of required or optional attribute types of an entry.

An example of the relationship between structural, abstract, and auxiliary object classes is the schema entry shown in Figure 16 on page 161. The schema entry specifies top, subEntry, subSchema, and **ibmSubschema** as object classes. The object classes form the following hierarchy:

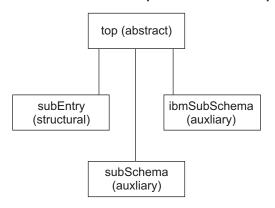


Figure 17. Object class hierarchy example

In this example, the **subEntry** object class is the base structural object class.

The OBSOLETE keyword indicates that the object class cannot be used to define entries in the directory. When an object class is made obsolete, new entries specifying the obsoleted object class cannot be added to the directory and existing entries cannot be modified unless the obsolete object class is removed from the entries' object class list. When the obsolete object class is removed from the entry, any attributes in the entry that are associated only with that object class must also be removed. These changes must be made through the same modify operation. If an obsolete object class is specified in a superior hierarchy for a new entry, then attempts to add the entry to the LDAP directory will fail.

LDAP syntaxes

Each attribute type definition includes the LDAP syntax which applies to the values for the attribute. The LDAP syntax defines the set of characters which are allowed when entering data into the directory.

The z/OS LDAP server is shipped with predefined supported syntaxes. See Table 26 on page 168 and Table 27 on page 168 for the list of syntaxes supported by the z/OS LDAP server. The set of syntaxes cannot be changed, added to, or deleted by users.

Matching rules

Matching rules allow entries to be selected from the database based on the evaluation of the matching rule assertion. Matching rule assertions are propositions which may evaluate to true, false, or undefined concerning the presence of the attribute value or values in an entry.

The z/OS LDAP server is shipped with predefined supported matching rules. See Table 28 on page 169 for the list of matching rules supported by the z/OS LDAP server. The set of matching rules cannot be changed, added to, obsoleted, or deleted by users.

Schema attribute syntax

The attributes which are used in the schema entry use specific character representations in their values. These character representations are described in Table 25. The terms shown in this table will be used in the schema attribute definitions in the next section.

Table 25. Character representations

Term	Definition
noidlen	Represented as:
	numericoid{length}
	where <i>length</i> is a numeric string representing the maximum length of values of this attribute type.
	Example:
	1.3.6.1.4.1.1466.115.121.1.7{5}
	Implementation note: The z/OS LDAP server allows values to be any length, regardless of the specification of a length in the attribute type definition. User installations that want to limit the length of values need to handle this during data input.
numericoid	A dotted decimal string.
	Example:
	2.5.13.72

Table 25. Character representations (continued)

Definition
A single object identifier. This may be specified either as a name or as a numeric object identifier.
Examples:
name 2.5.4.41
A list of object identifiers specified as names or numeric object identifiers separated by dollar signs (\$) within parentheses.
Example:
(cn \$ sn \$ postaladdress \$ 2.5.4.6)
Either an <i>oid</i> or <i>oidlist</i> .
A quoted description shown as 'descr' for one and as ('descr' 'descr') for more than one. The description (descr) must have an alphabetic character as the first character, followed by any combination of alphabetic or numeric characters, the dash character (-), or the semicolon character (;). Each value must be in single quotes (').
If there is more than one value, they must be enclosed in parentheses.
Examples:
'x121address' ('cn' 'commonName') 'userCertificate;binary'
Note: Although the LDAP V3 protocol does not support an underscore character (_) as a valid character in a <i>descr</i> , the z/OS LDAP server allows the use of an underscore character to facilitate data migration. This use should be minimized whenever possible and may not be supported by other servers.
A quoted descriptive string shown as 'dstring'. The descriptive string (dstring) can have any number of UTF-8 characters.
Example:
'This is an example of a quoted descriptive string.'

LDAP schema attributes

The five attributes used to define an LDAP schema are discussed below. For these schema attributes, the numericoid must be the first item in the definition. All other keywords and values may be in any order.

LDAP syntaxes

The set of syntaxes which are supported by the z/OS LDAP server cannot be modified, added to, or deleted by users. The descriptive material included here is for information only.

The format of the LDAP syntaxes attribute in a dynamic schema is:

ldapSyntaxes: (numericoid [DESC qdstring])

numericoid

The unique, assigned numeric object identifier.

DESC *qdstring*

Text description of the LDAP syntax

Note: LDAP syntaxes do not have a textual name. They are identified only by the numeric object identifier.

Following is an example of the definition of an LDAP syntax:

ldapSyntaxes: (1.3.6.1.4.1.1466.115.121.1.7 DESC 'Boolean')

The LDAP syntaxes supported by the z/OS LDAP server fall into two categories. The first set, as shown in Table 26, would be used when defining attribute types that are used for directory data.

Table 26. Supported LDAP syntaxes - general use

Numeric object identifier	Description	Valid values
1.3.6.1.4.1.1466.115.121.1.5	Binary	Binary data
1.3.6.1.4.1.1466.115.121.1.7	Boolean	TRUE, FALSE
1.3.6.1.4.1.1466.115.121.1.15	Directory String	UTF-8 characters
1.3.6.1.4.1.1466.115.121.1.12	Distinguished Name	Sequence of attribute type and value pairs
1.3.6.1.4.1.1466.115.121.1.24 Note: The effective time zone for the LDAP server is assumed when calculating GMT from local time.	Generalized Time	yyyymmddhhmmss.ffffff (local time) yyyymmddhhmmss.ffffffZ (GMT) yyyymmddhhmmss.ffffff-diff (local time with a differential to GMT, where diff is hhmm)
1.3.6.1.4.1.1466.115.121.1.26	IA5 String	IA5 characters (commonly known as 7-bit ASCII)
1.3.6.1.4.1.1466.115.121.1.27	Integer	+/- 62 digit integer
1.3.6.1.4.1.1466.115.121.1.38	Object Identifier	Name or numeric object identifier
1.3.6.1.4.1.1466.115.121.1.40	Octet String	Octet data
1.3.6.1.4.1.1466.115.121.1.50	Telephone Number	printable string (alphabetic, decimal, ", (,), +, ,, -, ., /, :, ?, and space)
1.3.6.1.4.1.1466.115.121.1.53	UTC Time	See Generalized Time above for details

The second set of syntaxes defined by the z/OS LDAP server are used in the definition of the LDAP schema. These would not typically be used in user schema attribute type definitions. They are listed here for reference.

Table 27. Supported LDAP syntaxes - server use

Numeric object identifier	Description
1.3.6.1.4.1.1466.115.121.1.3	Attribute Type Description
1.3.6.1.4.1.1466.115.121.1.16	DIT Content Rule Description
1.3.6.1.4.1.1466.115.121.1.17	DIT Structure Rule Description
1.3.18.0.2.8.1	IBM Attribute Type Description
1.3.18.0.2.8.3	IBM Entry UUID Description
1.3.6.1.4.1.1466.115.121.1.54	LDAP Syntax Description
1.3.6.1.4.1.1466.115.121.1.30	Matching Rule Description
1.3.6.1.4.1.1466.115.121.1.31	Matching Rule Use Description
1.3.6.1.4.1.1466.115.121.1.35	Name Form Description
1.3.6.1.4.1.1466.115.121.1.37	Object Class Description
1.3.6.1.4.1.1466.115.121.1.58	Substring Assertion

Matching rules

The set of matching rules which are supported by the z/OS LDAP server cannot be modified, added to, obsoleted, or deleted by users. The descriptive material included here is for information only.

The format of the matching rules attribute in a dynamic schema is:

matchingRules: (numericoid [NAME qdescrs] [DESC qdstring] [OBSOLETE] SYNTAX numericoid)

numericoid

The unique, assigned numeric object identifier.

NAME *qdescrs*

The name by which this matching rule is known.

DESC *qdstring*

Text description of the matching rule.

OBSOLETE

Indicates that the matching rule is obsolete.

SYNTAX numericoid

Specifies the numeric object identifier of the syntax for this matching rule.

Following is an example of the definition of a matching rule:

matchingRules: (2.5.13.5 NAME 'caseExactMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

The matching rules supported by the z/OS LDAP server is a fixed set as listed in the following table.

Table 28. Supported matching rules

Numeric object identifier	Name	Syntax
2.5.13.13	booleanMatch	1.3.6.1.4.1.1466.115.121.1.7 (Boolean)
1.3.6.1.4.1.1466.109.114.1	caseExactIA5Match	1.3.6.1.4.1.1466.115.121.1.26 (IA5 String)
2.5.13.5	caseExactMatch	1.3.6.1.4.1.1466.115.121.1.15 (Directory String)
2.5.13.6	caseExactOrderingMatch	1.3.6.1.4.1.1466.115.121.1.15 (Directory String)
2.5.13.7	caseExactSubstringsMatch	1.3.6.1.4.1.1466.115.121.1.58 (Substring Assertion)
1.3.6.1.4.1.1466.109.114.2	caseIgnoreIA5Match	1.3.6.1.4.1.1466.115.121.1.26 (IA5 String)
2.5.13.2	caseIgnoreMatch	1.3.6.1.4.1.1466.115.121.1.15 (Directory String)
2.5.13.3	caseIgnoreOrderingMatch	1.3.6.1.4.1.1466.115.121.1.15 (Directory String)
2.5.13.4	caseIgnoreSubstringsMatch	1.3.6.1.4.1.1466.115.121.1.15 (Directory String)
2.5.13.1	distinguishedNameMatch	1.3.6.1.4.1.1466.115.121.1.12 (Distinguished Name)
1.3.18.0.2.4.405	distinguishedNameOrderingMatch	1.3.6.1.4.1.1466.115.121.1.12 (Distinguished Name)
2.5.13.27	generalizedTimeMatch	1.3.6.1.4.1.1466.115.121.1.24 (Generalized Time)
2.5.13.28	generalizedTimeOrderingMatch	1.3.6.1.4.1.1466.115.121.1.24 (Generalized Time)

Table 28. Supported matching rules (continued)

Numeric object identifier	Name	Syntax
1.3.18.0.2.22.2	IBM-EntryUUIDMatch	1.3.18.0.2.8.3 (IBM Entry UUID)
2.5.13.29	integerFirstComponentMatch	1.3.6.1.4.1.1466.115.121.1.27 (Integer)
2.5.13.14	integerMatch	1.3.6.1.4.1.1466.115.121.1.27 (Integer)
2.5.13.0	objectIdentifierMatch	1.3.6.1.4.1.1466.115.121.1.38 (Object Identifier)
2.5.13.30	objectIdentifierFirstComponentMatch	1.3.6.1.4.1.1466.115.121.1.38 (Object Identifier)
2.5.13.17	octetStringMatch	1.3.6.1.4.1.1466.115.121.1.40 (Octet String)
2.5.13.20	telephoneNumberMatch	1.3.6.1.4.1.1466.115.121.1.50 (Telephone Number)
2.5.13.21	telephoneNumberSubstringsMatch	1.3.6.1.4.1.1466.115.121.1.58 (Substring Assertion)
2.5.13.25	utcTimeMatch	1.3.6.1.4.1.1466.115.121.1.53 (UTC Time)

Attribute types

The format of the attribute types attribute in a dynamic schema is:

attributeTypes: (numericoid [NAME qdescrs] [DESC qdstring] [OBSOLETE] [SUP oid] [EQUALITY oid][ORDERING oid] [SUBSTR oid] [SYNTAX noidlen] [SINGLE-VALUE] [NO-USER-MODIFICATION] [USAGE attributeUsage])

numericoid

The unique, assigned numeric object identifier.

NAME *qdescrs*

The name and alias names by which this attribute type is known. This is also known as the object identifier. The first name in the list is used as the base name and the other names are referred to as alias names. It is suggested the shortest name be listed first. If a name is not specified, the numeric object identifier is used to refer to the attribute type.

DESC *qdstring*

Text description of the attribute type.

OBSOLETE

Indicates that the attribute type is obsolete.

SUP oid

Specifies the superior attribute type. When a superior attribute type is defined, the **EQUALITY**, ORDERING, SUBSTR, and SYNTAX values may be inherited from the superior attribute type. The referenced superior attribute type must also be defined in the schema. When the SYNTAX, EQUALITY, ORDERING, or SUBSTR values are not specified for an attribute type, the attribute type hierarchy will be used to determine these values. The **SYNTAX** must be specified on the attribute type or through inheritance.

EQUALITY oid

Specifies the object identifier of the matching rule which is used to determine the equality of values.

ORDERING oid

Specifies the object identifier of the matching rule which is used to determine the order of values.

SUBSTR oid

Specifies the object identifier of the matching rule which is used to determine substring matches of values.

SYNTAX noidlen

The syntax defines the format of the data stored for this attribute type. It is specified using the numeric object identifier of the LDAP syntax and, optionally, the maximum length of data stored for this attribute type.

Implementation note: The z/OS LDAP server allows values to be any length, regardless of the specification of a length in the attribute type definition. User installations that want to manage the lengths of values need to handle this when values are put into the directory.

SINGLE-VALUE

Limits entries to only one value for this attribute type.

NO-USER-MODIFICATION

When specified, users may not modify values of this attribute type.

USAGE attributeUsage

Specify userApplications for attributeUsage. If USAGE is not specified, the default is userApplications.

The directoryOperation, distributedOperation, and DSAOperation keywords are reserved for LDAP server usage and are used to define the attribute types which comprise the schema entry itself. Attribute types defined by users must specify userApplications for usage or allow usage to default to userApplications.

Following are examples of the definition of attribute types:

```
attributeTypes: ( 2.5.4.6 NAME 'c' SUP name SINGLE-VALUE )
attributeTypes: ( 2.5.4.41 NAME 'name' EQUALITY caseIgnoreMatch SUBSTR
   caseIgnoreSubstringsMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{32768} )
```

IBM attribute types

The format of the IBM attribute types attribute in a dynamic schema is:

```
ibmAttributeTypes: ( numericoid [ACCESS-CLASS ibmAccessClass] )
```

numericoid

The unique, assigned numeric object identifier of the associated attribute type.

ACCESS-CLASS ibmAccessClass

The level of sensitivity of the data values for this attribute type. The acceptable values are **normal**, sensitive, and critical. See "Attribute access classes" on page 251 for definition of these values.

The IBMAttributeTypes schema element is an extension of the attributeTypes schema element. If the attributeTypes value is not defined, then the corresponding IBMAttributeTypes value cannot be defined.

Some schema elements are shipped with ACCESS-CLASS set to restricted or system. These values are reserved for LDAP server use. Other IBM LDAP servers may also specify DBNAME and LENGTH keywords and values. These keywords are not used by the z/OS LDAP server and do not need to be specified when creating schemas. If they are specified in a schema used by the z/OS LDAP server, they will be ignored.

Following is an example of the definition of an IBM attribute type:

```
ibmAttributeTypes: (2.5.4.6 ACCESS-CLASS normal)
```

Object classes

The format of the object classes attribute in a dynamic schema is:

```
objectClasses: ( numericoid [NAME gdescrs] [DESC gdstring]
   [OBSOLETE] [SUP oids] [ABSTRACT|STRUCTURAL |AUXILIARY] {[MUST oids] [MAY oids]} )
```

numericoid

The unique, assigned numeric object identifier.

NAME qdescrs

The name and alias names by which this object class is known. This is also known as the object identifier. The first name in the list is used as the base name. If name is not specified, the numeric object identifier is used to refer to the object class.

DESC qdstring

Text description of the object class.

OBSOLETE

Indicates that the object class is obsolete.

SUP oids

List of one or more superior object classes. When a superior object class is defined, entries specifying the object class must adhere to the superset of MUST and MAY values. The supersets of MUST and MAY values include all MUST and MAY values specified in the object class definition and all MUST and MAY values specified in the object class's superior hierarchy. When an attribute type is specified as a MUST in an object class in the hierarchy and a MAY in another object class in the hierarchy, the attribute type is treated as a MUST. Referenced superior object classes must be defined in the schema.

ABSTRACT | STRUCTURAL | AUXILIARY

Indicates the type of object class. **STRUCTURAL** is the default.

MUST oids

List of one or more mandatory attribute types. Attribute types which are mandatory must be specified when adding or modifying a directory entry.

MAY oids

List of one or more optional attribute types. Attribute types which are optional may be specified when adding or modifying a directory entry.

The extensibleObject object class is an AUXILIARY object class which allows an entry to optionally hold any attribute type. The extensibleObject object class is supported by the z/OS LDAP server. This allows any attribute type that is known by the schema to be specified in an entry which includes extensibleObject in its list of object classes.

The **top** object class is an abstract object class used as a superclass of all structural object classes.

Following is an example of the definition of an object class:

```
objectClasses: ( 2.5.6.0 NAME 'top' ABSTRACT MUST objectclass )
objectClasses: ( 2.5.6.6 NAME 'person' SUP top STRUCTURAL MUST ( cn $ sn )
MAY (userpassword $ telephonenumber $ seealso $ description ) ) objectClasses: (5.6.7.8 NAME 'company' SUP top MUST (department $ telephoneNumber ) MAY (postalAddress $ street ) ) objectClasses: (1.2.3.4 NAME 'companyPerson' SUP (company $ person ) )
```

Defining new schema elements

You can define new schema elements for use by applications that you develop to use the directory. You can add new object classes and attribute types to the schema defined to the TDBM backend in the server. To define a new object class or attribute type, create an LDIF file containing the new schema information, and perform an LDAP modify operation on the schema entry for the TDBM backend. Object classes and attribute types must be defined using the formats described in the previous section, and must include unique numeric object identifiers and names. Ensuring that the numeric object identifier and names are unique is essential to the correct operation of the directory when using your newly defined schema elements.

Numeric object identifiers (OIDs) are strings of numbers, separated by periods. OID "ranges" or "arcs" are allocated by naming authorities. If you are going to define new schema elements, you should obtain an

"OID arc" from a naming authority. One such location to get an "OID arc" assigned is managed by Internet Assigned Numbers Authority (IANA) and, at the time of this writing, was found at:

```
http://www.iana.org
```

Select the "Application Forms" link and then the "Private Enterprise Number" link to apply for a Private Enterprise number.

Once you have obtained an "OID arc" you can begin assigning OIDs to object classes and attribute types that you define.

For the example below, assume that we have been assigned OID arc 1.3.18.0.2.1000.100. (Note: Do not use this OID arc for defining your own schema elements. This arc is assigned to IBM for its use.) The following example adds a new object class that refers to two new attribute types. As you can see, the object class and attribute types can be added to the schema using a single LDAP modify operation. The changes to the schema are represented in LDIF mode input below:

```
dn: cn=schema, o=Your Company
changetype: modify
add: attributetypes
attributetypes: ( 1.3.18.0.2.1000.100.4.1 NAME 'YourCompanyDeptNo'
  DESC 'A users department number.'
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 EQUALITY caseIgnoreMatch
  USAGE userApplications
attributetypes: ( 1.3.18.0.2.1000.100.4.2 NAME 'YourCompanyEmployeeID'
   DESC 'A user employee ID.'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 EQUALITY caseIgnoreMatch
  USAGE userApplications
add: objectclasses
objectclasses: (1.3.18.0.2.1000.100.6.1 NAME 'YourCompanyPerson'
  DESC 'Attached to inetOrgPerson to add more attributes.'
  SUP top
  AUXILIARY
  MAY ( YourCompanyDeptNo $ YourCompanyEmployeeID )
)
```

The example above assumes that the suffix for the TDBM backend is o=Your Company. When you define your own schema, you should set the name of the schema (the "dn:" portion of the LDIF) to the name of the schema entry in the TDBM backend.

This short description has described how to update the schema in the TDBM backend with new schema elements. Defining new schema elements is a complex undertaking and is described in other publications.

Updating the schema

Attention

Updating the schema, if not done properly, can result in being unable to access data. Read this section thoroughly to avoid this situation.

When the z/OS LDAP server is first started, a minimum schema is initialized. This minimum schema defines the schema elements which are required by the z/OS LDAP server to read a schema entry. The minimum schema is shown in Appendix B, "Minimum schema for TDBM" on page 393 in LDAP search results format. This schema is stored as a directory entry with dn: cn=schema, <suffix>. After the z/OS LDAP server is first started, the schema attribute types and object classes needed for the directory should be added to the schema entry using the LDAP modify function. For example, to add the schema contained in the schema.user.ldif file, copy the file to a working directory, update the suffix in the file, and specify the file in the **Idapmodify** command as follows:

ldapmodify -h ldaphost -p ldapport -D adminDN -w passwd -f /home/myuser/schema.user.ldif

See"Setting up the schema for TDBM - new users" on page 159 for more information.

See z/OS: Security Server LDAP Client Programming for information on the LDAP utilities.

The operations that can be performed include adding, replacing, or deleting any object class or attribute type that is not part of the minimum schema definition required by the LDAP server.

The modifications (additions, changes, and deletions) specified by the LDAP modify function are applied to the schema entry which is retrieved from the database by the server. The resulting schema entry becomes the active schema and is used by the TDBM backend to verify that directory changes adhere to it. This schema entry is then stored back into the directory database.

Updates to the schema must be performed such that the schema fully resolves. This includes:

- All attribute types referred to in object classes must exist in the schema.
- · All superior attribute types or object classes must exist.
- Only the syntaxes and matching rules supported by TDBM may be specified in attribute type definitions.
- All attribute types referred to in IBM attribute type definitions must also be defined as attribute types.

The schema may also be changed when using the Idif2tdbm utility to load entries into the database. See "Idif2tdbm program" on page 127 for more information.

Data entries in the directory must adhere to the active schema. If the schema is changed such that an object class has a new must attribute type, then whenever an entry which is of that object class is added or modified, the new attribute type must be contained in the entry.

The schema entry is required and, therefore, attempts to delete the schema entry will fail.

The TDBM backend will manage dynamic schema changes relative to the sysplex (multiple LDAP servers operating against the same database and multiple threads operating against the database). Changes to schema affect all z/OS LDAP servers operating against the same DB2 tables in multi-server mode.

Note: If you delete schema elements from the active schema when information in the directory still exists which refers to those schema elements, you will lose access to that information. To avoid this situation, modify the schema elements you wish to "delete" by marking them as OBSOLETE rather than deleting their definitions from the schema entry. In this way, no new entries can be created using these schema elements and the existing entries which do use these schema elements will still be accessible. Existing entries that use OBSOLETE schema elements must be modified to use only non-OBSOLETE schema elements during the next modification of the entry in order for the modification to succeed.

Analyzing schema errors

Following is some information about the possible cause of some schema errors that may be encountered when updating schema:

· For enhanced readability, type:value pairs in LDIF files may be split across multiple lines. The indicator to LDIF that the subsequent lines are continuations is that the first character on the subsequent line is a space. This character is ignored by parsers and it is assumed that the next character immediately follows the previous line. Therefore, if a space is needed between the last value on one line and the first value on the subsequent line, a second space needs to exist on the subsequent LDIF line. Various reason codes related to unrecognized values may be issued.

- Attempts to delete or obsolete values that are in the minimum schema will fail. Changes to the minimum schema are not allowed and are ignored by the server.
- The IBM attribute type schema attribute is an extension to the associated attribute type in the schema. If the schema contains an IBM attribute type value for which an attribute type value is not defined, the schema update will fail. For example,

```
ibmAttributeTypes: ( 1.2.3.4 ACCESS-CLASS normal )
cannot be specified unless
attributeTypes: ( 1.2.3.4 Name 'sample' ... )
```

- While the UTC Time syntax is supported, usage of the Generalized Time syntax is recommended. For UTC Time syntax, year values between 70 and 99 assume 1970 to 1999 and values between 00 and
- When searching attribute type values of GMT or UTC Time syntax, use GMT syntax in the search filter rather than local time. All time values are stored in the data store as GMT times.

Retrieving the schema

69 assume 2000 to 2069.

is also defined.

The following sections describe how you can display the schema entry and also find the subschemaSubentry DN.

Displaying the schema entry

The following command shows how to search for the schema entry:

```
ldapsearch -h ldaphost -p ldapport -s base -b "cn=schema,<suffix>" "objectclass=subschema"
```

Replace < suffix> with the DN of a TDBM or SDBM suffix from the configuration file depending on which schema you want to publish.

Immediately after the server is started for the first time, this command produces the results shown in Appendix B, "Minimum schema for TDBM" on page 393 when the <suffix> specified is managed by TDBM. After the schema has been updated by the administrator, the search results will show the full schema as the union of the minimum schema and the added schema elements. If the <suffix> specified in Idapsearch is managed by SDBM, the SDBM schema shown in Appendix C, "SDBM schema" on page 397is displayed. The SDBM schema is not modifiable.

The search results when searching the schema entry will show the distinguished name of the schema exactly as entered in the search request. For example:

```
ldapsearch -h ldaphost -p ldapport -s base -b "cn=SCHEMA,o=Acme Company,c=US" "objectclass=subschema"
```

produces:

```
cn=SCHEMA,o=Acme Company,c=US
cn=schema
subtreespecification=NULL
objectclass=TOP
objectclass=SUBSCHEMA
objectclass=SUBENTRY
objectclass=IBMSUBSCHEMA
attributetypes = ( 2.5.4.3 NAME ( 'cn' 'commonName' ) SUP name )
ibmattributetypes = ( 2.5.4.3 ACCESS-CLASS normal )
objectclasses = ( 2.5.6.0 NAME 'top' ABSTRACT MUST objectclass )
```

```
Idapsyntaxes = ( 1.3.6.1.4.1.1466.115.121.1.15 DESC 'directory string' )
matchingrules = ( 2.5.13.5 NAME 'caseExactMatch' SYNTAX '1.3.6.1.4.1.1466.115.121.1.15 )
```

Finding the subschemaSubentry DN

The subschemaSubentry attribute in each directory entry contains the DN of the schema that was active when the entry was added to the directory. To find the value of the subschemaSubentry attribute, specify subschemaSubentry as an attribute on an LDAP search of the entry.

```
ldapsearch -h ldaphost -p ldapport -s base -b "o=Acme Company, c=UK" "objectclass=*" subschemasubentry
o=Acme Company, c=UK
subschemasubentry= cn=SCHEMA, o=ACME COMPANY, C=UK
```

Migrating the schema from RDBM to TDBM

z/OS LDAP users who are migrating directories from the RDBM backend to the TDBM backend may also need to migrate the directory schema. The schema would need to be migrated if there had been any changes made to the previously shipped schema files or schema definitions that were added to the schema. The schema2ldif tool assists with this conversion. This should be done by the person who understands the schema in place for the directory.

IBM supplied schema

The schema.user.ldif and schema.IBM.ldif files are the LDAP V3 protocol versions of the schema.system.at, schema.system.oc, schema.user.at, schema.user.oc, schema.lBM.at, and schema.IBM.oc files that were used in previous releases of the z/OS LDAP server for RDBM. All schema elements contained in the listed *.at and *.oc files are contained either in schema.user.ldif or schema.IBM.Idif. The schema.user.Idif and schema.IBM.Idif files are used with the TDBM backend. The *.ldif files are different from the *.at and *.oc files in the following ways:

- · LDIF format versus configuration file format.
- Defined in terms of the LDAP Version 3 protocol (IETF RFC 2252 format).
- Use of new syntaxes and separate specification of matching rules.
- The syntaxes supported by RDBM in previous LDAP servers (bin, ces, cis, dn, tel) have been converted to a combination of a syntax and matching rule. These changes are consistent with the LDAP Version 3 protocol definitions of the attribute types.
- · Addition of new attribute types and object classes.
- Use of superior attribute types and object classes in definitions.
- For example, the definition of **cn** specifies the attribute type **name** as a superior.
- The schema.IBM.Idif file requires that the schema.user.Idif file be loaded first.
- If there were no customer-specific changes or additions made to the previously shipped schema, then the schema.user.ldif and schema.IBM.ldif files that are shipped with the z/OS LDAP server should be used to load the schema when migrating the directory.

User-defined schema

To minimize the migration effort, it is recommended that customer-specific defined attribute types and object classes be placed in separate .at and .oc files prior to migration. The output of this would be a separate LDIF file containing the customer-specific schema. This would then be used in conjunction with the z/OS LDAP *.ldif schema files located in /usr/lpp/ldap/etc.

The schema2ldif utility

The schema2ldif utility takes RDBM and SDBM attribute type and object class files and converts them to the standard LDIF format needed for TDBM.

The schema2ldif utility is Java based, so a jar and script file are provided with the z/OS LDAP server. The jar file, named servertools.jar, and the script file, named schema2ldif, are located in the /usr/lpp/ldap/sbin directory. The script file provides an easy manner to run the utility.

This utility can only be run from Unix System Services (USS).

Before running the schema2ldif utility, the PATH and CLASSPATH environment variables in the schema2ldif script file should be uncommented and updated if necessary. The export statements listed below are commented out in the schema2ldif script file.

- export CLASSPATH=/usr/lpp/ldap/sbin/servertools.jar:\$CLASSPATH
- export PATH=/usr/lpp/java/J1.3/bin:\$PATH

Check with your system administrator to find the proper location of Java on your system and adjust the **PATH** and **CLASSPATH** appropriately in the **schema2ldif** script file.

If there are include statements in any of your files, those additional files will be processed after the input schema files specifically entered on the command line are parsed and formatted in LDIF format.

Note: The schema2ldif utility will add SUP top to all object class definitions it encounters.

Format

```
schema2ldif -s schemaInputFile1 [schemaInputFile2...] -d ldifOutputFile -1 logFile [-f overrideFile]
  [-o OIDFile] [-v] [-?]
```

Parameters

- -? Specifies help text for the **schema2ldif** utility.
- -s schemalnputFile1 [schemalnputFile2...]

Specifies the name of one or more input files to be converted. These files may be attribute type, object class, or slapd.conf files from previous releases of OS/390 LDAP. If attribute types and object classes are mixed together in one input schema file, they must be separated into two different files in order for schema2ldif to properly parse the file to LDIF.

Following is an example of an attribute type file:

```
attribute
                                ces
                                                    50
                                     name
                                                        normal
                                                  11
attribute
         acceleratorCapabilities
                                cis
                                     accelCap
                                                        normal
attribute accessHint
                                dn
                                                  1000 normal
                                     accessHint
attribute accountHint
                               dn accountHint 1000
                                                        normal
attribute accountService
                               dn accountService 1000
attribute AccountSuffix
                               dn AccountSuffix 1000
                                                        normal
attribute actionPending
                               cis actionPending 5 normal
```

Following is an example of an object class file:

```
objectclass
             cacheObject
     requires
             objectClass
     allows
objectclass
             cimBIOSelement
     requires
             objectClass
     allows
             primaryBIOS
```

Following is a portion of an OS/390 Release 8 LDAP configuration file that has server parameters and includes for other files which are the actual attribute and object class files. In this case, the oc.conf and at.conf files will be opened and parsed to the output LDIF file.

```
#referral ldap://ldap.itd.umich.edu
include oc.conf
include at.file
sslKeyRingFile key.kdb
sslKeyRingFilePW key.kdb
sslCipherSpecs 12288
```

-d IdifOutputFile

Specifies the name of the LDIF output file. If there is already a file name the same as the output file specified on the command line, the original file name is renamed to:

```
ldifOutPutFile.1
```

All of the output from either the attribute types or the object classes input files are routed to one output file concatenated together.

-I loaFile

Specifies the name of the log file. After each invocation of the schema2ldif utility, a log file is created. If this file is empty, schema2ldif completed successfully without errors. Otherwise, errors are printed to this file and to standard output.

-f overrideFile

Specifies an optional override file. The purpose of this file is to modify attribute types and object classes from previous releases of the OS/390 or z/OS LDAP server that are encountered in the input schema files. It is strongly recommended that the system.override file, shipped with the z/OS LDAP server and located in /usr/lpp/ldap/etc, be used.

The format of the override file is:

```
{attributeType Name|objectClass Name} {ADD|DELETE|MODIFY} [new_attributeType|new_objectClass]
{attributeType_Name | objectClass_Name}
```

Name of the attribute type or object class that is being operated on.

ADD | DELETE | MODIFY

Operation that the schema2ldif utility will perform to the attribute type or object class.

```
new attributeType | new objectClass
```

Definition of the attribute type or object class as defined in previous versions of the LDAP server. This field is used when the operation is **ADD** or **MODIFY**.

DELETE

Does not put the specified attribute type or object class definition into the output LDIF file.

ADD

Adds new attribute types or object classes to the output LDIF schema file. For an ADD operation of an attribute type, the entire definition must be specified on one line. To add an object class, there must be four arguments specified on the first line: the object class name, ADD, the word objectclass, and the new object class name. This is followed on subsequent lines with the keyword requires or allows, and the lists of required or allowed attribute type

Attribute types or object classes will only be added to the output LDIF file if a respective attribute type or object class input file was encountered either on the command line or through an **include** statement.

MODIFY

Modifies an existing attribute type or object class. For a MODIFY of an attribute type, the entire modified attribute type definition must be specified on one line. For a MODIFY of an object class, the entire modified object class definition must be specified.

If any of the attribute types or object classes in the first column of any line denoting a **DELETE** or MODIFY are encountered in any of the input schema files, then the appropriate action is taken.

Notes:

- 1. Comments are denoted by a line starting with a # sign.
- Blank lines are ignored.
- When the schema2ldif utility encounters an erroneous line, the line is ignored and a warning messages is printed to standard out and the log file.

Following is an example of a system.override file:

```
#system.override file
top DELETE
faxnumber DELETE
telephone MODIFY attribute tphone ces telephone 32 critical
referral MODIFY objectclass refl
   requires
     objectClass
computer ADD attribute computer ces computer 50 normal
location new ADD objectclass location new
   requires
     objectClass, ou
   allows
     businessCategory
```

-o OIDFile

Specifies an optional numeric object identifier file. The purpose of this file is to provide the numeric object identifiers for the attribute types and object classes which were specified in the input *.at, *.oc, or *.conf files.

The format of this file is shown here:

2.5.4.41 name country 2.5.6.2 telephoneNumber 2.5.4.20 facsimileTelephoneNumber 2.5.4.23

The attribute type or object class name is the first field and the numeric object identifier is the second field.

- 1. Comments in the OID file are denoted by a line starting with a # sign. Any line that contains more than two fields is an error. The utility will display an "ATTENTION" message on the console and ignore the erroneous line.
- If an attribute type or object class name-numeric object identifier pair is specified twice in the OID file, the second numeric object identifier is used in the output LDIF file.

The OS/390 LDAP server provides a default OID file called name-noid.list, which is located in the /usr/lpp/ldap/etc directory.

If you have additional attribute types or object classes in your schema files, be certain that any additional name-numeric object identifier pairs you add do not conflict with the object identifiers that are in the name-noid.list file.

-v Specifies verbose mode, where all debugging messages are printed to the log file. If the -v option is specified, the log file will contain messages that state what the schema2ldif utility has done to the input schema file or files. In this file, all of the attribute types and object classes that were formatted to LDIF format are denoted along with the input schema file name. Messages denoting that the optional input files (that is, any OID or override files) have been parsed are also written to the log file.

Examples

Following is an example attribute type file named at.in:

```
attribute name
                                 cis
                                        name
                                                     50
                                                          normal
attribute acceleratorCapabilities cis
                                        acce1Cap
                                                    11
                                                          normal
attribute accessHint
                                 dn
                                        accessHint 1000
                                                          norma1
```

Figure 18. Example attribute type file (at.in)

Following is an example object class file named **oc.in**:

```
objectclass
              cacheObject
        requires
                objectClass
        allows
                tt1
objectclass
             cimBIOSelement
        requires
                objectClass
        allows
                primaryBIOS
```

Figure 19. Example object class file (oc.in)

The examples of schema2ldif that follow use these input files.

1. To convert at.in and oc.in to LDIF, use the following command:

```
schema2ldif -s at.in oc.in -d output.ldif -l log.out
```

This schema2ldif utility produces the following in the output.ldif file.

```
dn: cn=schema,<namingcontext>
changetype:modify
add:x
attributetypes: ( name-at-oid NAME 'name' SYNTAX '1.3.6.1.4.1.1466.115.121.1.15' EQUALITY caseIgnoreMatch )
ibmAttributeTypes: ( name-at-oid ACCESS-CLASS normal )
attributeTypes: ( acceleratorCapabilities-at-oid NAME 'acceleratorCapabilities' SYNTAX '1.3.6.1.4.1.1466.115.121.1.15'
 EQUALITY caseIgnoreMatch )
ibmAttributeTypes: ( acceleratorCapabilities-at-oid ACCESS-CLASS normal )
attributeTypes: ( accessHint-at-oid NAME 'accessHint' SYNTAX '1.3.6.1.4.1.1466.115.121.1.12' EQUALTIY distinguishedNameMatch )
ibmAttributeTypes: ( accessHint-at-oid ACCESS-CLASS normal )
objectclasses: (cacheObject-oc-oid NAME 'cacheObject' SUP TOP MUST (objectClass) MAY (ttl))
objectclasses: ( cimBIOSelement-oc-oid NAME 'cimBIOSelement' SUP TOP MUST ( objectClass ) MAY ( primaryBIOS ) )
```

Note that this generated LDIF file cannot be used directly with the z/OS LDAP server because the numeric object identifiers need to be added manually.

2. This example uses both at.in and oc.in as schema input files, but specifies an OID file named oid.in shown below.

```
2.5.4.41
name
accessHint
                1.3.18.0.2.4.292
cacheObject
                1.3.6.1.4.1.250.3.18
```

The following **schema2ldif** command:

```
schema2ldif -s at.in oc.in -d mixed.out -l log.out -o oid.in
```

produces the following output in **mixed.out**:

```
dn: cn=schema,<namingcontext>
changetype:modify
add:x
```

```
attributeTypes: ( 2.5.4.41 NAME 'name' SYNTAX '1.3.6.1.4.1.1466.115.121.1.15' EQUALITY caseIgnoreMatch ) ibmAttributeTypes: ( 2.5.4.41 ACCESS-CLASS normal ) attributeTypes: ( acceleratorCapabilities-at-oid NAME 'acceleratorCapabilities' SYNTAX '1.3.6.1.4.1.1466.115.121.1.15' EQUALITY caseIgnoreMatch ) ibmAttributeTypes: ( acceleratorCapabilities-at-oid ACCESS-CLASS normal ) attributeTypes: ( 1.3.18.0.2.4.292 NAME 'accessHint' SYNTAX '1.3.6.1.4.1.1466.115.121.1.12' EQUALTIY distinguishedNameMatch) ibmattributeTypes: ( 1.3.18.0.2.4.292 ACCESS-CLASS normal ) objectclasses: ( 1.3.6.1.4.1.250.3.18 NAME 'cacheObject' SUP TOP MUST ( objectClass ) MAY ( ttl ) ) objectclasses: ( cimBIOSelement-oc-oid NAME 'cimBIOSelement' SUP TOP MUST ( objectClass ) MAY ( primaryBIOS ) )
```

3. This example specifies a slapd.conf file (shown below) as input to the schema2ldif utility.

```
include at.in
include oc.in
sslKeyRingFile key.kdb
sslKeyRingFilePW none
sslCipherSpecs 12288
```

The following **schema2ldif** command:

```
schema2ldif -s slapd.conf -d mixed.out -l mixed.log -v
```

produces the following output in mixed.out:

The log file **mixed.log** will have the following information since the **-v** setting has been turned on:

```
PROCESSING input file slapd.conf
PROCESSING include file at.in
FORMATTING attribute type name
FORMATTING attribute type acceleratorCapabilities
FORMATTING attribute type accessHint
PROCESSING include file oc.in
FORMATTING objectclass cacheObject
FORMATTING objectclass cimBIOSelement
```

4. This example uses an override file. The **slapd.conf** from Example 3, the **at.in** file from Figure 18 on page 180, and the **oc.in** file from Figure 19 on page 180 files are used in this example. Following is the **system.override** file also used in this example:

```
# system.override file
name DELETE
cacheObject DELETE
owner ADD attribute owner dn owner 1000 normal
accessHint MODIFY attribute hint dn hint 1000 normal
cimBIOSelement MODIFY objectclass BIOSelement
    requires
    objectclass
allows
    primaryBIOS
```

The following **schema2ldif** command:

```
schema2ldif -s slapd.conf -d mixed.out -l mixed.log -v -f system.override
```

produces the following output in **mixed.out**:

```
ibmAttributeTypes: ( acceleratorCapabilities-at-oid ACCESS-CLASS normal )
attributeTypes: ( hint-at-oid NAME 'hint' SYNTAX '1.3.6.1.4.1.1466.115.121.1.12' EQUALITY distinguishedNameMatch )
ibmAttributeTypes: ( hint-at-oid ACCESS-CLASS normal )
objectclasses: ( BIOSelement-oc-oid NAME 'BIOSelement' SUP TOP MAY ( primaryBIOS ) )
```

The **mixed.log** file produces the following output:

```
ADDING: new attribute (owner) from system.override file to ADD table
ADDING: name from override file system.override to DELETE table
ADDING: cacheObject from override file system.override file to DELETE table
ADDING: accessHint from override file system.override to MODIFY table
ADDING: cimBIOSelement from override file system.override to MODIFY table
PROCESSING input file slapd.conf
PROCESSING include file at.in
SKIPPING attribute type name. This attribute type was included as a DELETE in the override file.
FORMATTING attribute type owner
FORMATTING attribute type acceleratorCapabilities
USING modified attribute type hint from override file
FORMATTING hint
PROCESSING include file oc.in
SKIPPING objectclass type cacheObject. This objectclass was included as a DELETE in the override file.
USING modified object class BIOSelement from override file
FORMATTING BIOSelement
```

Note

Some changes to the resulting schema LDIF file may be required depending on the level of the schema files used during migration.

- · Check for multiple definitions of the cRLDistributionPoint, pilotDSA, and qualityLabelledData object classes. Remove the duplicate definition from the LDIF file.
- · Check for the existence of the container object class. If it does not exist in the schema file, add it as defined below:

```
objectclasses: ( 1.3.18.0.2.6.28 Name 'Container' SUP TOP MUST ( cn ) )
```

Add SUP person to the definition of the organizationalPerson object class.

Chapter 17. Modify DN Operations

The Modify DN Operation allows a client to change the leftmost (least significant) component of the name of an entry in the directory, or to move a subtree of entries to a new location in the directory. This chapter explains the function of the Modify DN operation and the options supported to influence the scope and duration of the operation. In addition, it instructs on the techniques necessary to achieve certain forms of directory renames and movement, and it advises on issues which may result in unintentional or unwanted results.

Modify DN Operation Syntax

The z/OS implementation of the Modify DN operation supports all required and optional parameters described for the operation in RFC 2251. Specifically, these parameters are required:

- **entryDN**: This is the Distinguished Name (DN) of the entry whose name will be changed. This entry may or may not have subordinate entries. This parameter may not be a zero-length string.
- newRdn: The Relative Distinguished Name (RDN) that will form the leftmost component of the new
 name of the entry. This parameter may not be a zero-length string. If the intent of the Modify DN
 operation is to move the target entry to a new superior without changing its RDN, the old RDN value
 must be supplied in the newRdn parameter.
- **deleteoldrdn:** A boolean parameter that controls whether the old RDN attribute values are to be retained attributes of the entry or whether they will be deleted from the entry.

The following parameter to the Modify DN operation is optional:

• **newSuperior:** The Distinguished Name (DN) of the entry which will become the immediate superior of the renamed entry (identified by the **entryDN** parameter). If this parameter is present, it may consist of a zero-length string or a non-zero-length string. See "Modify DN operations related to suffix DNs" on page 193 for more information on the use of a zero-length string for this parameter. A zero-length string value for this parameter ("") will signify that the new superior entry is the root DN.

This operation also supports optional values, or controls, to influence the behavior of the operation. Two controls are supported (see Appendix G, "Supported server controls" on page 431):

- IBMModifyDNTimelimitControl: This control causes the Modify DN operation to be abandoned if its duration exceeds the time limit represented by the control value expressed in seconds. No changes are made if the operation is abandoned. If the server's time limit is less than the time limit requested in the control for this operation, the server's time limit will take precedence. See "Configuration file options" on page 53 for more information on the servers time limit. This control is honored even if it is set by the admin DN for the server. When this control is present, it will not be propagated to the replica servers. (See "Modify DN operations and replication" on page 200 for more information about replication of Modify DN operations.)
- IBMModifyDNRealignDNAttributesControl: This control causes the server to search for all attributes whose attribute type is based on a DN syntax (designated by OID 1.3.6.1.4.1.1466.115.121.1.12) and whose values match any of the old DN values being renamed as part of the Modify DN operation, and to modify the old DN values to reflect the corresponding renamed DN attribute values. This includes modifications to two other attribute types which have constructed DN-type attribute values (those whose attribute syntax is not distinguished name but which may be used to store DN values). They are aclEntry and ownership entryOwner attributes. Updates to constructed DN types will be limited to these two attributes defined by the LDAP Directory Server. No changes will be made to any user constructed types.

This control is an all-or-none operation in which the server attempts to realign all appropriately-matched DN attribute values in the TDBM backend. Users cannot limit the scope of values which should be realigned. If a failure arises during the realignment operation, it realigns none of the values, and the Modify DN operation fails. No changes are made if the operation is abandoned. It should be noted that

even if the control is designated as non-critical, the server will still try to honor the intent of the control and if this attempt fails, the entire Modify DN operation will fail.

When IBMModifyDNRealignDNAttributesControl is present on a request to a master server on which replication of Modify DN operations is enabled, it will be propagated to the replica servers. (See "Modify DN operations and replication" on page 200 for more information about replication of Modify DN operations.)

A few simple examples of the use of the Modify DN operation follow. Each request will be expressed in the format of the ModifyDNRequest defined in RFC 2251, as well as in the corresponding invocation command for the z/OS client utility program Idapmodrdn. Refer to the z/OS: Security Server LDAP Client Programming for more information on the **Idapmodrdn** utility.

Example 1: Simple Modify DN of leaf node

```
ModifyDNRequest ::= {
entry
                cn=Kevin Heard, o=Athletics, o=Human Resources, o=Deltawing, c=AU
newrdn
                cn=Kevin T. Heard
deleteoldrdn TRUE
```

ldapmodrdn -h ldaphost -p ldapport -D binddn -w passwd -r cn=Kevin Heard, o=Athletics, o=Human Resources, o=Deltawing, c=AU" "cn=Kevin T. Heard"

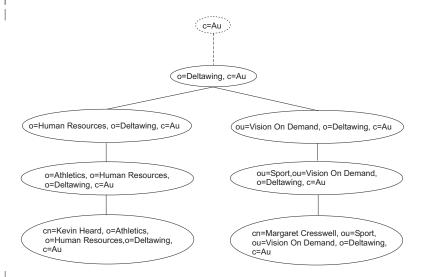


Figure 20. Before Modify DN operation

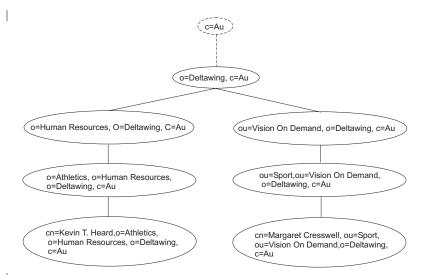


Figure 21. After Modify DN operation

Note: The **-r** parameter specifies that the old RDN attribute value (cn=Kevin Heard) will be deleted from the target entry after this operation.

Example 2: Simple Modify DN of non-leaf node

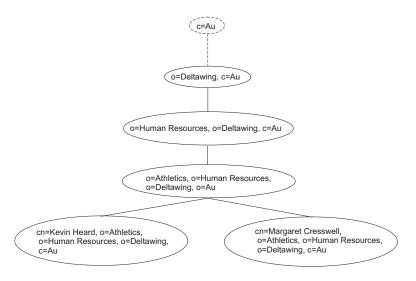


Figure 22. Before Modify DN operation

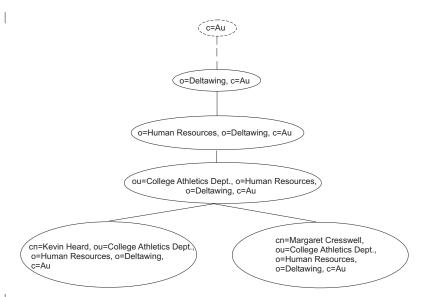


Figure 23. After Modify DN operation

Note: The absence of the -r parameter specifies that the old RDN attribute value (o=Athletics) will be preserved in the target entry after this operation.

Example 3: Modify DN of non-leaf node with relocation (newSuperior)

```
ModifyDNRequest ::= {
                o=Athletics, o=Human Resources, o=Deltawing, c=AU
entry
                o=Adult Athletics
newrdn
deleteoldrdn
                FALSE,
newSuperior
                ou=Sport, ou=Vision On Demand, o=Deltawing, o=AU
ldapmodrdn -h ldaphost -p ldapport -D binddn -w passwd -s "ou=Sport, ou=Vision On
Demand, o=Deltawing, c=AU" "o=Athletics,o=Human Resources, o=Deltawing, c=AU" "o=Adult Athletics"
```

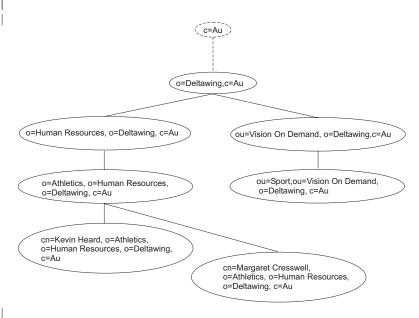


Figure 24. Before Modify DN operation

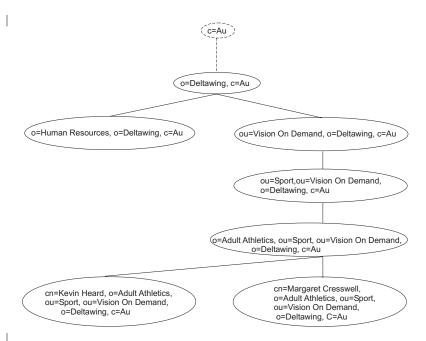


Figure 25. After Modify DN operation

Note: The absence of the **-r** parameter specifies that the old RDN attribute value (o=Athletics) will be preserved in the target entry after this operation. The target entry and descendants in its subtree will be relocated in the directory hierarchy.

Considerations in the use of Modify DN operations

As this operation has the potential to significantly change directory data and how it can be accessed, it is important that the user fully understand the data before using the Modify DN operation. Specifically, the user needs to know that:

- The ability of this operation to move directory subtrees has the potential for affecting many entries in the directory in a single operation.
- Certain options may result in modification of additional directory entries which are outside the scope of the directory subtrees being moved. This chapter will explain and give examples of how that can occur.
- Because the changes performed to the directory as a result of the operation are committed as a single transaction (or reversed if an error occurs), it may result in a long-running transaction, which may reduce concurrency of other LDAP operations targeted for the same directory entries. See "Concurrency considerations between Modify DN operations and other LDAP operations" on page 188 for more information.
- The scope of the changes may result in unanticipated effects in the directory and may affect user access to these entries. See "Access control changes" on page 191 for more information.
- There are limitations to which directory entries are eligible for the Modify DN operation. See "Eligibility of entries for rename" on page 188 for more information.
- In case the directory needs to be returned to a state prior to a Modify DN operation, the directory should be backed up by using the tdbm2ldif utility program or by using DB2 utilities to generate a DB2 image copy of the underlying tablespaces. See Chapter 11, "Running and using the LDAP backend utilities" on page 125 for more information about the tdbm2ldif utility program, and DB2: Universal Database for OS/390 Master Index for more information about the DB2 image copy. In addition to backing up the directory contents, activity logging should be enabled before nontrivial changes are made to the directory. See "Activity log and console listing of Modify DN operations" on page 203 for more information.
- There are considerations if the data to be modified by this operation is being replicated. See "Modify DN operations and replication" on page 200 for more information.

Eligibility of entries for rename

Entries in the directory which are targeted to be renamed in a single Modify DN operation are subject to these constraints:

- 1. All entries to be renamed must be located in the same TDBM backend instance targeted by the Modify DN operation. The Modify DN operation with newSuperior option will move subtree entries within the same TDBM backend instance, and will not permit movement of subtree entries from one backend instance to another. The operation will only relocate entries to subtrees conforming to the same schema. The entry to be renamed must exist in the backend, and the new DN for the entry must not already exist in the backend.
- 2. Referral entries may be renamed as part of a Modify DN operation. If a referral entry is renamed as part of a Modify DN operation, its corresponding entry in another backend must be manually updated to reflect the name changes; no automatic updates are propagated to those backends from the target backend. Referrals which exist in other directory servers which refer to any of the entries whose DNs were modified in the local directory by a Modify DN operation will need to be manually updated to reflect the changes; no automatic updates are propagated to those servers from the local one.
- 3. Schema entries may not be renamed explicitly by the Modify DN operation. However, if a TDBM suffix entry is renamed, and the schema entry DN included that suffix, then the schema entry will be automatically renamed to reflect the new DN.
- 4. Entries renamed by a Modify DN operation must conform to the schema rules for the backend in which the operation occurs. As such, the RDN attribute type must be consistent with the schema rules for the object classes of the entry: a Modify DN operation will fail if the attribute type of newRdn is not in the MUST or MAY list for the entry's object classes.
- 5. When IBMModifyDNRealignDNAttributesControl is present on a Modify DN request, the operation modifies occurrences of the renamed DN if the syntax for the attribute containing the matching DN value is distinguished name. If a subtree of entries is moved and the entries are thereby renamed, then the operation modifies all occurrences of all renamed DNs in the. In addition, if a matching DN value has an attribute syntax which is one of the two constructed attribute types, aclEntry or entryOwner, the matching values will be replaced with their renamed values. Any user constructed attribute types which contain DNs whose values match those being modified by the Modify DN operation will be left unchanged.
- 6. A Modify DN operation can succeed even if the newRdn value already exists as an attribute value in the entry, as long as that value is not already the sole RDN attribute value. For example, suppose an entry with DN of dept=AAA,ou=mydivision,o=MyCompany,c=us is to be renamed with the newRdn sector=northeast, and that the entry already contains the multi-valued attribute sector which currently contains attribute values of northeast and northcentral. This rename will succeed, and the value northeast will appear only once for the attribute type **sector**.
- 7. Entries may be renamed only if all access control requirements are satisfied for the bound user, as determined by the effective ACL and ownership permissions for those entries and attributes. See "Access control and ownership" on page 189 for detailed explanation and examples of this effect.

Concurrency considerations between Modify DN operations and other LDAP operations

The ability of the Modify DN operation to rename non-leaf nodes in the directory (which causes all entries which are hierarchical subordinates of the target entry to be renamed) and the ability to move directory subtrees has the potential for affecting many entries in the directory in a single operation. In addition, use of IBMModifyDNRealignDNAttributesControl with this operation may result in modification of additional directory entries which are outside the scope of the directory subtrees being renamed or moved. Changes to all entries affected by the operation are committed at the same time. While modified entries are awaiting the transaction commit point, database locks are held which prevent other concurrent operations from sharing and modifying the data. If many entries undergo modification with this operation, it may result in a long-running transaction which has potential for reducing concurrency of other operations targeted for the same directory entries. Although the LDAP server is capable of processing concurrent LDAP operations

targeted at a given TDBM backend instance while the Modify DN operation is in progress, the extent to which such concurrency is possible will depend on what data in the directory may be needed and locked by the competing operations. If the number of entries being affected by the Modify DN operation is large, the underlying DB2 locking mechanism may elect to escalate locking levels, which would result in more entries being excluded from use by other concurrent operations than those which will be modified by the Modify DN operation. Thus, users should be mindful of this when submitting Modify DN operations to the server. If such an operation is expected to affect large numbers of entries in the directory, it may be advisable to submit such a request during a low-activity period when demand for the same resources by multiple concurrent operations is relatively low.

Access control and ownership

For all entries being renamed, the caller must have **w**(rite) permissions for the attribute values that will have to change in all affected entries. In addition, if the *newSuperior* parameter is present on the Modify DN request, the caller must have permissions of **object:a** on the *newSuperior* entry and **object:d** on the target entry at the top of the subtree of entries being moved. If the caller lacks one or more of these permissions, the operation is denied. No access control checking is done against any of the target entry's subordinates even though their DN is changed. It should be noted that if the caller is an effective owner of any of the entries being renamed, the permissions are automatically satisfied for those entries.

In addition, if the **IBMModifyDNRealignDNAttributesControl** accompanies a Modify DN request, then the bound DN must have **w**(rite) permission to all of the attributes that are changed as a result of realignment of the DN values.

Example:

Assume our sample directory contains the following entry which will be the target of a Modify DN operation, and which contains explicit ACL information:

```
dn: o=Athletics, o=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU
aclEntry: access-id: cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd.,
o=Deltawing,c=AU:normal:rswc:sensitive:rsc:object:d

(other attributes not shown)
```

The directory also contains an entry with DN ou=Production, ou=Vision On Demand,o=Deltawing, c=AU which will be the newSuperior of the Modify DN operation. This entry inherits the following ACL information (propagated from a superior entry):

```
aclEntry: access-id: dn: cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU:normal:rwsc:sensitive:rsc:object:a
```

In addition, there are several entries containing attributes of DN syntax. For this example, assume that these attribute types and their respective attribute access classes are as follows:

```
attribute: reportingOrganizationaccess-class: sensitiveattribute: workingOrganizationaccess-class: normal
```

The LDIF format representation of the entries containing *reportingOrganization* or *workingOrganization* attributes are:

```
dn: cn=Lisa Fare, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU cn: Lisa Fare objectclass: organizationalPerson objectclass: person objectclass: TOP aclEntry: access-id: cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing,c=AU:normal:rsc:sensitive:rs sn: Fare
```

```
title: Occupational Health and Safety Administrator
telephonenumber: (07) 635 1432
manager: cn=John Gardner, ou=Human Resources Group, ou=Deltawing InfoSystems,
o=Deltawing, c=AU
secretary: cn=Ian Campbell, o=Deltawing, c=AU
reportingOrganization: o=Athletics, o=Human Resources, ou=Delta Home Media Ltd.,
o=Deltawing, c=AU
dn: cn=Laurie Wood, ou=Human Resources Group, ou=Deltawing Automotive Ltd., o=Deltawing, c=AU
cn: Laurie Wood
objectclass: organizationalPerson
objectclass: person
objectclass: TOP
aclEntry: access-id: cn=Mark Crawford, ou=Human Resources, ou=Delta Home
Media Ltd., o=Deltawing,c=AU:normal:rswc:sensitive:rsw
telephonenumber: (03) 9335 2114
title: Pay Officer
workingOrganization: o=Athletics, o=Human Resources, ou=Delta Home Media Ltd.,
o=Deltawing, c=AU
```

Relocating an entry

User "cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing,c=AU" submits the following Modify DN operation request to the server to relocate the target entry:

ldapmodrdn -h ldaphost -p ldapport -D "cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing,c=AU" -w passwd -s "ou=Production, ou=Vision On Demand, o=Deltawing, c=AU" "o=Athletics, o=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU" "o=Athletics Division"

The -s parameter specifying newSuperior is present on this operation request, so in addition to the access permissions needed for all Modify DN operations (w on affected attributes), the user also needs object:d on the target entry and object:a on the newSuperior entry. The bound user is in the aclEntry for the target entry as well as in the aclEntry for the newSuperior entry, and has all required access permissions (can write attributes and delete the target entry, and can add objects under the newSuperior entry), so the operation is permitted.

Relocating an entry with DN realignment requested

Now suppose the same user submits a Modify DN operation request to the server to relocate the same target entry under the same newSuperior entry, but with the addition of the control requesting realignment of DN attribute values (**-a** parameter):

ldapmodrdn -h *ldaphost* -p *ldappart* -D "cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing,c=AU" -w passwd -a -s "ou=Production, ou=Vision On Demand, o=Deltawing, c=AU" "o=Athletics, o=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU" "o=Athletics Division"

In addition to the permissions required on the previous example, this operation now requires additional permissions to be checked on entries containing values which qualify for realignment. The DN being modified ("o=Athletics, o=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU") is found in DN-syntax attributes of two entries: The entry with DN "cn=Laurie Wood, ou=Human Resources Group, ou=Deltawing Automotive Ltd., o=Deltawing, c=AU" contains this value in the workingOrganization attribute, and the entry with DN "cn=Lisa Fare, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU" contains this value in the reportingOrganization attribute.

The bound user is in the aclEntry for "cn=Laurie Wood, ou=Human Resources Group, ou=Deltawing Automotive Ltd., o=Deltawing, c=AU". The workingOrganization attribute is in the access-class of normal, and the bound user is granted w access to this class of attributes, so the realignment of the DN value would be permitted in this entry.

The bound user is also in the **aclEntry** for "cn=Lisa Fare, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing, c=AU". The **reportingOrganization** attribute is in the access-class of sensitive, and the bound user is granted only **rs** permissions on **sensitive** attributes in the entry, so the realignment of this value would be denied. Even though the bound user had adequate permissions to perform the relocation of the target entry and had adequate permissions to perform realignment of the DN value in one of the two entries containing a matching DN, the operation would fail because the bound user does not have the necessary permissions on everything needed to complete the operation.

Access control changes

If a Modify DN operation is accompanied by the *newSuperior* parameter changes in effective ACLs and in effective ownership of the relocated entries may result. Regardless of the effective ACLs which applied to the moved subtree in its old location, the moved subtree now inherits any propagating ACLs applying to the *newSuperior* entry. As a consequence, entries to which a user had access before the request may no longer be accessible by that user, and entries to which access was denied for a given user before the request may now be accessible by that user.

Explicit ACLs in the entry or subtree override propagating ACLs. All explicit ACLs which were in the moved subtree at its original location move along with the entries.

When renaming a DN, it is possible that ACLs and entryOwners containing the renamed DN will be modified. Therefore, prior to such a move or rename users should carefully consider how ownership and accessibility to entries protected by these attributes may change after the move, and what ACL and ownership changes may be desired, if any.

The following is an example of how a Modify DN operation might affect access controls:

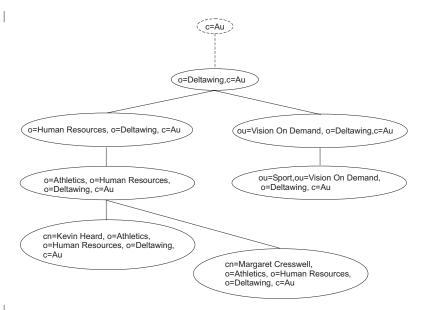


Figure 26. Before Modify Dn operation

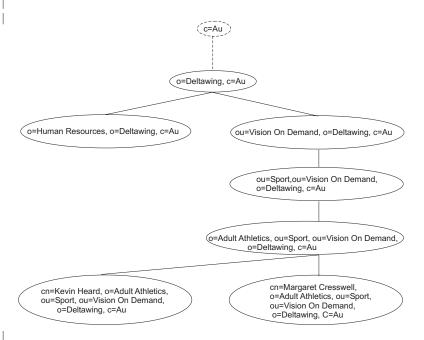


Figure 27. After Modify DN operation

Assume that the entry with DN o=Human Resources, o=Deltawing, c=AU has an explicit propagating ACL containing the following aclEntry:

aclEntry: access-id: cn=Mark Edmondson, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing,c=au:normal:rwcs:sensitive:rwcs:critical:rws:object:d

Also, assume that the entry with DN ou=Sport, ou=Vision On Demand, o=Deltawing, c=AU has an explicit propagating ACL containing the following aclEntry:

aclEntry: access-id: cn=Mark Edmondson, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing,c=au:normal:rws:sensitive:r:critical:r:object:a

If the user bound as DN cn=Mark Edmondson, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing,c=AU performs the example Modify DN operation, there are at least two consequences which should be noted:

- While this DN previously had rwcs permissions on sensitive attributes in the entry o=Athletics, o=Human Resources, o=Deltawing, c=AU and rws permissions on critical attributes in the same entry, this DN now has only r access on both sensitive and critical attributes in the entry after the relocation. It might be expected that a given DN will have the same accessibility to specific entries and data in the directory after a Modify DN operation as it had to those entries and data before the operation, but this example demonstrates that such an expectation is not valid.
- If, after completion of the Modify DN operation, the bound user decides that they wish to return the moved entry (and its subordinates) back to their original location in the directory hierarchy, this will not be possible with the access controls currently in place. The bound DN has only **object:d** permission on the old superior node ("o=Human Resources, o=Deltawing, c=AU") where **object:a** is needed to effect the move of an entry or subtree under the superior node, and the bound DN has only **object:a** permission on the moved entry ("o=Adult Athletics, ou=Sport, ou=Vision On Demand, O=Deltawing, c=AU") where **object:d** is needed to move the entry. Thus, while it may be expected that a given DN can reverse a Modify DN operation under all circumstances, this example demonstrates that such an expectation is not valid.

Ownership changes

When the *newSuperior* parameter accompanies the Modify DN request, any entries in a relocated subtree which had explicit owners prior to the relocation will preserve that explicit ownership after the relocation has been performed. Any entries in the relocated subtree which inherited ownership prior to relocation will continue to inherit ownership following relocation. If the owning entry prior to relocation was a node superior to the relocated entry, the owning entry will be the new superior entry. If the owning entry was an entry within the relocated subtree, the owning entry is preserved following the relocation.

Any entries in the relocated subtree which propagated ownership to subordinates prior to relocation continue to propagate ownership to subordinates after the relocation.

Refer to the example in the preceding section "Access control changes" on page 191.

Assume that the entry with DN o=Human Resources, o=Deltawing, c=AU has an explicit propagating owner of cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd.,o=Deltawing,c=AU.

Also, assume that the entry with DN ou=Sport, ou=Vision On Demand, o=Deltawing, c=AU has an explicit propagating owner of cn=Neville McAuliffe, ou=Human Resources Group, ou=Deltawing Infosystems, o=Deltawing, c=AU.

Before the Modify DN operation, the effective owner of the renamed entry is cn=Mark Crawford, ou=Human Resources, ou=Delta Home Media Ltd., o=Deltawing,c=AU; after completion of the operation, the effective owner of the renamed entry is now cn=Neville McAuliffe, ou=Human Resources Group, ou=Deltawing Infosystems, o=Deltawing,c=AU. Thus, the act of relocating an entry may change the effective owner of that entry and of its subordinates.

Modify DN operations related to suffix DNs

The Modify DN operation may be used to modify the DNs of any and all entries in a given TDBM backend. In addition to renaming leaf entries (directory entries with no subordinate entries) and mid-hierarchy entries (directory entries which have both superior entries and subordinate entries), suffix entries may also be renamed. Suffix entries may be renamed to become non-suffix entries and suffix entries may be renamed such that they continue to be suffix entries. In addition, non-suffix entries may be renamed to become suffix entries. This section provides example scenarios for rename operations which involve suffix entries. It summarizes constraints which have been adopted for the z/OS LDAP implementation which are not

defined in the protocol behavior prescribed by RFC 2251 for the Modify DN operation. Examples are provided on how various renaming scenarios may be accomplished, and factors to be considered when performing these operations are discussed.

Scenario constraints

When using a Modify DN operation to move or rename a suffix DN, numerous scenarios should be considered. A Modify DN operation must be completely contained within a given instance of a TDBM backend. Both source entries and target entries must be in the same backend. Each of the sample renaming scenarios which follow requires that the entry must exist and the result of a rename must not yield a new entry which already exists in this backend instance. If either one or both of these assumptions are violated, the operation fails.

Several constraints will apply which are not defined by RFC 2251 in the description of the protocol behavior:

- 1. If an entry being renamed will become (or remain) a suffix, the new DN must be designated in the server's configuration file as a suffix for the backend, otherwise the operation will not be permitted.
- 2. The newRdn parameter of the Modify DN request must contain a non-null value, otherwise the operation request will be treated as an error.
- 3. If the *newSuperior* parameter is present, it may contain a zero-length string signifying that the newSuperior entry is the root DN.

In the directory hierarchy diagrams which follow, a circle outlined with a dashed line represents a component of a suffix DN. Circles containing gray fill represent DNs for which an entry exists in the directory.

Example scenarios

The following are example scenarios:

1. Rename a suffix RDN with no accompanying newSuperior, and the new DN remains a suffix after the rename is completed.

Example:

Suffixes defined in the server configuration file:

```
suffix: ou=End GPL, o=MyCompany, c=US
suffix: ou=Endicott, o=MyCompany, c=US
```

Rename operation is to rename suffix object entry ou=End GPL, o=MyCompany, c=US to suffix object entry ou=Endicott, o=MyCompany, c=US

The following figure is an example of this operation:

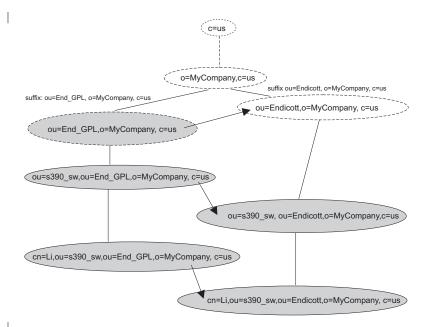


Figure 28. Suffix rename with no new superior

The new DN must be already designated as a suffix for this backend, otherwise this operation will fail.

The operation is performed the same as a rename of any other RDN in the directory

 a. Send Modify DN operation request with target=ou=End_GPL, o=MyCompany, c=US newRdn=ou=Endicott

This results in renaming ou=End_GPL, o=MyCompany, c=US to ou=Endicott, o=MyCompany, c=US and in renaming subordinate entries accordingly.

2. Rename of suffix DN with an accompanying newSuperior, and the new DN remains a suffix after the rename is completed. Example:

Suffix defined in the server configuration file: suffix: ou=Endicott, o=MyCompany, c=us

Rename operation is to rename suffix object entry ou=Endicott, o=MyCompany, c=us to suffix object entry o=MyCompany, c=us

The following figure shows an example of this operation:

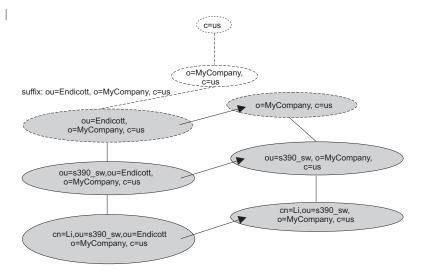


Figure 29. Suffix rename with new superior

This scenario, which involves renaming an existing suffix to an overlapping new suffix, must be performed in several steps, since the product does not permit designation in the server configuration file of overlapping suffixes for the same backend instance. The definition of overlapping suffixes is when two suffixes with differing numbers of naming components are equal to the extent of the shorter of the two suffixes. For example, ou=Endicott, o=MyCompany, c=US and o=MyCompany,c=US are considered to be overlapping suffixes, while ou=Endicott, o=MyCompany, c=US and ou=Raleigh, o=MyCompany, c=US are not considered to be overlapping suffixes.

This rename can be accomplished by having a temporary suffix pre-defined for the backend (for example, o=0urTemporarySuffix), renaming the target entry to become the temporary suffix, stopping the server and deleting the suffix ou=Endicott, o=MyCompany, c=us and adding the suffix o=MyCompany, c=us, and restarting the server. The temporary suffix would later be deleted from the list of suffixes for the backend.

a. Send a Modify DN operation request with target= ou=Endicott, o=MyCompany, c=us newRdn= o=OurTemporarySuffix newSuperior= "" (present in request with zero-length string)

This results in renaming ou=Endicott, o=MyCompany, c=us to o=OurTemporarySuffix. Note that the server treats newRdn as an error if it contains a zero-length string, but zero-length strings are permitted in the *newSuperior* argument to signify that the superior entry is the root DN.

b. Stop server, remove suffix ou=Endicott, o=MyCompany, c=us from the server configuration file, add suffix o=MyCompany, c=us, and restart server.

This results in adding the desired target suffix without a resulting conflict from overlapping suffixes.

c. Send a Modify DN operation request with target= o=OurTemporarySuffix newRdn= o=MyCompany newSuperior= c=us

This step results in renaming the temporary suffix o=0urTemporarySuffix to the desired suffix o=MyCompany, c=us, thereby accomplishing the rename from ou=Endicott, o=MyCompany, c=us to o=MyCompany, c=us. In the process, subordinate entries would be renamed accordingly.

3. This example shows the renaming of a suffix to another overlapping suffix higher in the directory hierarchy. A similar scenario could also be performed involving the rename of a suffix to another overlapping suffix, where the new name is a suffix lower in the directory hierarchy. Example:

Suffix defined in the server configuration file suffix: ou=Endicott, o=MyCompany, c=us

Rename operation is to rename suffix entry ou=Endicott, o=MyCompany, c=us to suffix object entry div=S390, ou=Endicott, o=MyCompany, c=us

The following figure shows an example of this operation:

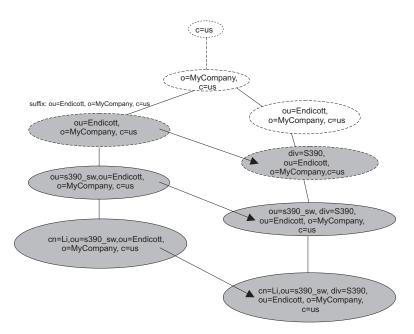


Figure 30. Overlapping suffix rename A

This rename can be accomplished by having a temporary suffix pre-defined for this backend in the server configuration file (for example, o=OurTemporarySuffix), renaming the target entry to become the temporary suffix, stopping the server and deleting the suffix ou=Endicott, o=MyCompany, c=us and adding the suffix div=S390, ou=Endicott, o=MyCompany, c=us, and restarting the server. The temporary suffix would later be deleted from the list of suffixes for the backend. This scenario would be done like so:

- a. Send a Modify DN operation request with target= ou=Endicott, o=MyCompany, c=us newRdn= o=OurTemporarySuffix newSuperior= "" (present in request with zero-length string)
- b. Stop server, remove suffix ou=Endicott, o=MyCompany, c=us, add suffix div=S390, ou=Endicott, o=MyCompany, c=us, and restart server.
- c. Send a Modify DN operation request with target= o=OurTemporarySuffix newRdn= div=S390 newSuperior= ou=Endicott, o=MyCompany, c=us

At this point, it should be noted that if these operational scenarios are to be replicated from a master server to one or more replica servers, there is a procedure which must be followed to permit this.

- a. Stop the replica server(s), add the temporary suffix (o=OurTemporarySuffix in our examples), restart the replica server(s).
- b. Stop the master server, perform the previous Steps 3a and 3b from the examples above. This will result in the intermediate rename to be performed on the master server and the results to be propagated to the replica server(s).

- c. Stop the replica server(s), delete the original suffix (ou=Endicott, o=MyCompany, c=us in both examples above), add the new suffix (o=MyCompany, c=us in the first example above, div=S390, ou=Endicott, o=MyCompany, c=us in the second example above), and restart the replica server(s).
- d. Perform the previous Step 3c from the examples above. This will result in the rename of entries to the final destination on the master server and in the results being propagated to the replica server(s).
- 4. Rename of suffix DN (some component other than RDN), and the new DN remains a suffix after the rename is completed. Example:

Suffixes defined in the server configuration file:

suffix: ou=Endicott, o=MyCompany, c=us suffix: ou=Endicott, o=MyCompany ny, c=us

Rename operation is to rename suffix entry object ou=Endicott, o=MyCompany, c=us to suffix entry object ou=Endicott, o=MyCompany ny, c=us

The following figure shows an example of this operation:

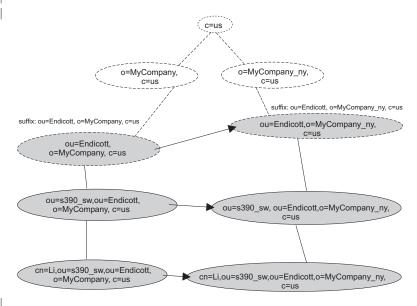


Figure 31. Overlapping suffix rename B

The new DN must be already designated as a suffix for this backend, otherwise this operation will fail. The operation is performed the same as a rename of any other DN in the directory. The product will permit the rename to occur in one step, even if an object entry for newSuperior does not already exist, since the newly-named object will become a suffix entry.

Send a Modify DN operation request with target= ou=Endicott, o=MyCompany, c=us newRdn= ou=Endicott newSuperior= o=MyCompany_ny, c=us

This results in renaming the DN from ou=Endicott, o=MyCompany, c=us to u=Endicott, o=MyCompany ny, c=us and in renaming subordinate entries accordingly.

5. Rename of suffix DN (including some component other than RDN), with an accompanying newSuperior, but the new DN is no longer a suffix Example:

Suffixes defined in the server configuration file: suffix: ou=End, o=MyCompany, c=us

suffix: ou=End,ou=MyCompany na, o=MyCompany, c=us

Rename operation is to rename suffix entry object ou=End, o=MyCompany, c=us to non-suffix entry object ou=GPL, ou=End, ou=MyCompany na, o=MyCompany, c=us

The following figure shows and example of this operation:

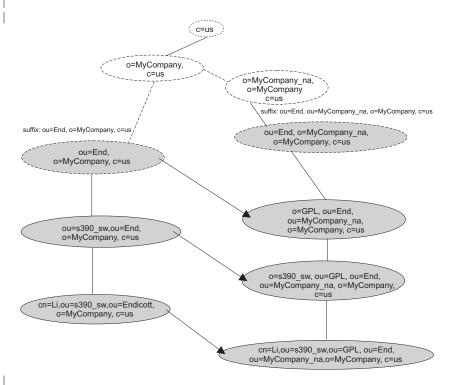


Figure 32. Suffix rename to non-suffix entry

The newSuperior entry object must already exist before this operation will be permitted.

a. Send a Modify DN operation request with

```
target= ou=End, o=MyCompany, c=us
newRdn= ou=GPL
newSuperior= ou=End, ou=MyCompany na, o=MyCompany,c=us
```

This results in renaming ou=End, o=MyCompany, c=us to ou=GPL, ou=End, ou=MyCompany_na, o=MyCompany, c=us and in renaming subordinate entries accordingly.

6. Rename of a non-suffix DN (including some component other than RDN), with an accompanying newSuperior, and the new DN is now a suffix Example:

Suffixes defined in the server configuration file:

```
suffix: ou=End, o=MyCompany, c=us suffix: o=Lotus, c=us
```

Rename operation is to rename non-suffix div=Lotus,ou=End, o=MyCompany, c=us to suffix o=Lotus, c=us

The following figure shows and example of this operation:

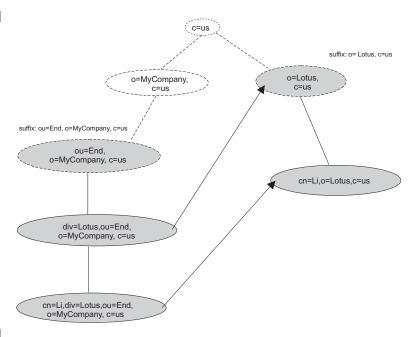


Figure 33. Rename non-suffix entry to suffix entry

The new DN must be already designated as a suffix for this backend, otherwise this operation will fail.

a. Send a Modify DN operation request with target= div=Lotus,ou=Endicott, o=MyCompany, c=us newRdn= o=Lotus newSuperior= c=us

This step results in renaming div=Lotus, ou=Endicott, o=MyCompany, c=us to o=Lotus, c=us and in renaming subordinate entries accordingly.

Modify DN operations and replication

Modify DN operations may be classified into two categories:

- 1. Simple Modify DN operations are those which rename a leaf node, and which are not accompanied by the newSuperior parameter or the IBMModifyDNRealignDNAttributesControl control or the IBMTimeLimitControl control.
- 2. Complex Modify DN operations are those which either rename a mid-tree (non-leaf) node, or which are accompanied by the newSuperior parameter, or which are accompanied by either the IBMModifyDNRealignDNAttributesControl control or the IBMTimeLimitControl control.

Simple Modify DN operations are always accepted by the master server, and are replicated if replicaObjects are present in the TDBM backend where a Modify DN operation is applied.

Complex Modify DN operations are only accepted by the master server, and are propagated to replica servers contingent on a determination that replica servers are compatible server versions. A compatible server version is one known to support for Modify DN operations all features and controls implemented by the z/OS Release 4 LDAP server including:

- the IBMModifyDNRealignDNAttributesControl control
- the IBMTimeLimitControl control
- the *newSuperior* parameter
- rename of non-leaf entries (complex Modify DN operations).

If one or more of these features or controls is not supported by a replica server, all complex Modify DN operations will be refused at the master server.

Periodic validation of compatible server versions in replica servers

Periodic checks are made of replica servers by the master server which are intended to increase the likelihood that complex Modify DN operations will be successfully replicated. Following is a description of the mechanisms used by master servers to do such checking.

The LDAP server must be able to establish a connection to each of the replica servers represented by replicaObjects in a given TDBM backends. When the connection is established to a given replica server, the master server determines if the replica server is at a compatible server version based on a query of the root DSE on that server. If a connection cannot be established to a replica server, it is assumed that the server does not provide the requisite support for replication of Modify DN operations, and complex Modify DN operations are refused at the master server. If a connection is established to a replica server and it is determined that the replica is not at a compatible server version, complex Modify DN operations are refused at the master server. Note that replication of simple Modify DN operations is always permitted, and such operations are always performed at the master server.

During operation of the master server, it may enable or disable processing of complex Modify DN operations, dependent on dynamically changing states of replica servers and of replicaObjects within the master server's TDBM backend. It is possible for the server to refuse complex Modify DN operations after having accepted them for some period of time, and it is possible for the server to accept complex Modify DN operations after having refused them for some period of time. Such a change can be triggered by several events. Each replication cycle tests connections to all replica servers defined by replicaObjects in the TDBM backend, and if a connection can no longer be established to at least one of the replica servers (even if it had been established to the same replica on the previous replication cycle), the master server begins refusing complex Modify DN operations. If all connections succeed but it is determined that one or more of the replica servers is not at a compatible server version (such as might happen, for example, when the replica server has been stopped when running one version of the LDAP server code and subsequently restarted using a different version of the LDAP server code), the master server begins refusing complex Modify DN operations. Only if connections may be established successfully to all replica servers and if they are determined to be running a compatible server version will the master server resume accepting complex Modify DN operations.

Other possible events which may influence whether the master server accepts or refuses complex Modify DN operations involve:

- the addition of new replicaObjects
- deletion of existing replicaObjects
- modification of existing replicaObjects in the TDBM backend.

Each of these causes the master server to temporarily suspend processing of complex Modify DN operations, until the check of replica servers at the start of the next replication cycle, at which point the replica server version levels will be used to determine whether the master server resumes accepting complex Modify DN operations.

Loss of replication synchronization due to incompatible replica server versions

The LDAP Server replication model runs periodically, rather than continuously, and the state of the replica is not checked until the start of each replication cycle. A complex Modify DN operation could be accepted or rejected based on inaccurate information about the state of a replica server between the start of two replication cycles. As a consequence, the replication process could stall and the synchronization between the master server and its replicas could be lost.

Attention

It is highly recommended that before starting a master server containing replica objects which may be the recipient of complex Modify DN operations, that the LDAP server administrator ensure that replica servers are each at a compatible server version level.

This will preclude any situations where replication will become stalled due to the presence of complex Modify DN operations which can not be propagated to all servers in the replica ring. To determine whether a replica server is at a compatible version level, submit a root DSE search to that server, similar to the following:

ldapsearch -h *ldaphost* -p *ldapport* -v 3 -s base -b "" objectclass=*

where *Idaphost* represents the hostname on which the replica server runs and *Idapport* is the port number on which the replica server is listening.

If the attribute type ibmDirectoryVersion is returned on the root DSE search, its value must equal z/OS V1R4 to be a compatible server version. If the ibmDirectoryVersion attribute is not returned on the root DSE search, or if the ibmDirectoryVersion attribute is returned on the root DSE search but its value is not z/OS V1R4, the server is an incompatible server version.

Loss of replication synchronization due to incompatible replica server versions - recovery

If at some point a master server accepts a complex Modify DN operation which can not be replicated, there are several means of recovering from this situation. The best method of recovering from this situation is to ensure that all replica servers are reachable from the master server, and that all replica servers are running at a compatible version level (this may entail stopping some replica servers and restarting them at a compatible version level). Once this state has been reached, gueued changes awaiting propagation to replica servers will drain from the queue at the master server and the replication process will resume normal operation.

An alternative is to delete the replicaObject from the master server corresponding to the replica server which is currently unreachable or which is running at an incompatible server level. Note that this will result in loss of synchronization with that replica server, and if one wishes to later restart the offending replica (such as, after it has been brought up to a compatible server version) it will be necessary to take a backup of the master server contents and restore those contents to the replica server before restarting it, to ensure the two directories are synchronized.

Replication of Modify DN operations, shared databases, and compatibility

The z/OS R4 LDAP Server is backward-compatible with earlier releases of the LDAP Server with respect to replication of Modify DN operations. If an earlier release of the LDAP Server stores replication change records for Modify DN operations in a TDBM backend which is subsequently used by a z/OS V1R4 LDAP Server, the V1R4 server is capable of replicating Modify DN operations which have not yet been propagated to the replica servers. This backward-compatibility provides for migration to the z/OS V1R4 LDAP Server from earlier server versions which support the TDBM backend.

It should be noted that once a z/OS V1R4 LDAP Server has performed Modify DN operations (of either simple or complex varieties) and stored replication change records in a given DB2-based TDBM backend, an earlier release of the z/OS LDAP Server should not be run against this backend. Replication change records for Modify DN operations written by the z/OS V1R4 LDAP Server are not recognized by earlier versions of the z/OS LDAP Server, and will cause failure of the replication process.

Activity log and console listing of Modify DN operations

The server will make use of the Activity Log (see "Activity Logging" on page 100) and the system console log to record the parameters of all successful Modify DN operations. This will include all values in the operation request, including control values and criticalities. By recording the details of the Modify DN operation request.

To record the Modify DN operation details in the Activity Log, the MSGS level of Activity Log operation must be enabled (see "Activity Logging" on page 100).

It should be noted that while the server is running, it may be necessary to perform a flush of the Activity Log (see "Activity Logging" on page 100) to view all messages for Modify DN operations performed to date.

Regardless of whether the Activity Log is configured to record information about Modify DN operations, the operation details (parameter values and control values) will be written to the system console log unconditionally for all successful complex Modify DN operations. If the customer determines this information is needed for diagnostic purposes within a limited amount of time after the directory data has been changed, it will still be available if archived copies of the system console log are maintained. Following is sample Activity Log output from Modify DN operations, with descriptions of the operation parameters and controls which generated the messages:

Example Modify DN operation submitted to the server:

```
ldapmodrdn -h ldaphost -p ldapport -D adminDN -w passwd -l 100 -r "o=01,o=Deltawing,c=AU" "o=01 Renamed"
```

This operation specifies a time limit of 100 seconds, renames the DN o=01,o=Deltawing,c=AU by modifying the RDN to o=01 Renamed, and deletes the old RDN value from the entry.

Following is the Activity Log output produced by successful completion of this operation:

```
Mon Feb 4 10:36:27 2002 GLD0215I End Successful Modify DN operation: dn =>
 o=01,o=Deltawing,c=AU; newrdn => o=01 Renamed; deleteoldrdn => TRUE
Mon Feb 4 10:36:27 2002 GLD0216I End Successful Modify DN operation: dn =>
o=01,o=Deltawing,c=AU; newSuperior => Not present in operation request.
Mon Feb 4 10:36:27 2002 GLD0217I End Successful Modify DN operation: dn =>
 o=01,o=Deltawing,c=AU; control type=> 1.3.18.0.2.10.10; control criticality => TRUE;
 control value (hexadecimal) => x'f1f0f000'
```

Example Modify DN operation submitted to the server:

```
ldapmodrdn -h ldaphost -p ldapport -D adminDN -w passwd -l 300 -a "ou=Sport,
ou=Production, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing, c=AU"
"ou=Renamed Sport"
```

This operation specifies a time limit of 300 seconds, requests that DN realignment be performed, and renames the DN ou=Sport, ou=Production, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing, c=AU by modifying the RDN to ou=Renamed Sport.

Following is the Activity Log output produced by successful completion of this operation:

```
Mon Feb 4 09:44:24 2002 GLD0215I End Successful Modify DN operation: dn => ou=Sport,
ou=Production, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing, c=AU;
newrdn => ou=Renamed Sport; deleteoldrdn => FALSE
Mon Feb 4 09:44:24 2002 GLD0216I End Successful Modify DN operation: dn => ou=Sport,
ou=Production, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing, c=AU;
newSuperior => Not present in operation request.
Mon Feb 4 09:44:24 2002 GLD0217I End Successful Modify DN operation: dn => ou=Sport,
ou=Production, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing, c=AU;
control type => 1.3.18.0.2.10.11; control criticality => TRUE; control value
(hexadecimal) => x'Not present in operation request.'
```

Mon Feb 4 09:44:24 2002 GLD0217I End Successful Modify DN operation: dn => ou=Sport. ou=Production, ou=Vision On Demand, ou=Delta Home Media Ltd., o=Deltawing, c=AU; control type => 1.3.18.0.2.10.10; control criticality => TRUE; control value (hexadecimal) => x'f3f0f000'

Example Modify DN operation submitted to the server:

ldapmodrdn -h ldaphost -p ldapport -D adminDN -w passwd -l 10500 -a -s "" "cn=Steve Lawlor, ou=English, ou=Languages, ou=Education Requirements Research, ou=Market Research, ou=Home Ed, ou=Delta Home Media Ltd., o=Deltawing,c=AU" "o=TopOfDir"

This operation specifies a time limit of 10500 seconds, requests that DN realignment be performed, renames the DN cn=Steve Lawlor, ou=English, ou=Languages, ou=Education Requirements Research, ou=Market Research, ou=Home Ed, ou=Delta Home Media Ltd., o=Deltawing,c=AU by modifying the RDN to o=TopOfDir, and relocates the entry to the root of the directory using the newSuperior parameter (-s "").

Following is the Activity Log output produced by successful completion of this operation:

Mon Feb 4 10:06:20 2002 GLD0215I End Successful Modify DN operation: dn => cn=Steve Lawlor, ou=English, ou=Languages, ou=Education Requirements Research, ou=Market Research, ou=Home Ed, ou=Delta Home Media Ltd., o=Deltawing, c=AU; newrdn => o=TopOfDir; deleteoldrdn => FALSE Mon Feb 4 10:06:20 2002 GLD0216I End Successful Modify DN operation: dn => cn=Steve Lawlor, ou=English, ou=Languages, ou=Education Requirements Research, ou=Market Research, ou=Home Ed, ou=Delta Home Media Ltd., o=Deltawing,c=AU; newSuperior => Mon Feb 4 10:06:20 2002 GLD0217I End Successful Modify DN operation: dn => cn=Steve ${\tt Lawlor,\ ou=English,\ ou=Languages,\ ou=Education\ Requirements\ Research,\ ou=Market}$ Research, ou=Home Ed, ou=Delta Home Media Ltd., o=Deltawing,c=AU; control type => 1.3.18.0.2.10.11; control criticality => TRUE; control value (hexadecimal) => x'Not present in operation request.' Mon Feb 4 10:06:20 2002 GLD0217I End Successful Modify DN operation: dn => cn=Steve Lawlor, ou=English, ou=Languages, ou=Education Requirements Research, ou=Market Research, ou=Home Ed, ou=Delta Home Media Ltd., o=Deltawing,c=AU; control type => 1.3.18.0.2.10.10; control criticality => TRUE; control value (hexadecimal) => x'f1f0f5f0f000'

Note: In the sample Activity Log output above, when the newSuperior parameter contains a zero-length string (""), the Activity Log message containing the newSuperior value for the operation is printed as a blank.

Chapter 18. Accessing RACF information

RACF provides definitions of users and groups, as well as access control for resources. The LDAP server can provide LDAP access to the user and group information stored in RACF.

Using SDBM, the RACF database backend of the LDAP server, you can:

- · Add new users and groups to RACF
- Add users to groups (connections)
- · Modify RACF information for users and groups
- Retrieve RACF information for users and groups
- · Delete users and groups from RACF
- Remove users from groups (connections)

The SDBM database of the LDAP server implements portions of the **adduser**, **addgroup**, **altuser**, **altgroup**, **deluser**, **delgroup**, **listuser**, **listgrp**, **connect**, **remove**, and **search** RACF commands. (See "SDBM operational behavior" on page 212 for more information.) An individual user has the same authority through SDBM as with normal RACF commands. The SDBM database of the LDAP server makes use of the R_Admin "run command" interface to accomplish its access to RACF data. As a result, this support is subject to the restrictions of the R_Admin interface. See *z/OS: Security Server RACF Callable Services* for more information regarding these restrictions. One restriction in particular affects return of search results. Refer to "RACF restriction on amount of output" on page 219 for more details.

The SDBM database allows for directory authentication (or bind) using the RACF user ID and password. The RACF user ID must have an OMVS segment defined and an OMVS UID present. The RACF user and group information that make up an identity can be used to establish access control on other LDAP directory entities. This expands use of the RACF identity to the rest of the LDAP-managed namespace. Note the following when using RACF access:

- An LDAP simple bind to a z/OS LDAP server using RACF access support but having a non-RACF security manager will succeed as long as the __passwd() call made by the LDAP server is successful. However, no group membership information will be available for the bound distinguished name if the security manager is not RACF.
- An LDAP simple bind made to a z/OS LDAP server using RACF access support continues to provide a
 successful or unsuccessful LDAP return code. In addition, if the LDAP return code being returned is
 LDAP_INVALID_CREDENTIALS, additional information is provided in the "message" portion of the
 LDAP result. The additional information is an LDAP-unique reason code and reason code text in the
 following format:

Rnnnnn text

The following errno values returned by __passwd() will have an LDAP reason code defined for them:

Table 29. The errno values returned by _passwd()

errno value	Reason	Text
EMVSEXPIRE	R000100	The password has expired.
EMVSPASSWORD	R000101	The new password is not valid.
EMVSSAFEXTRERR	R000102	The userid has been revoked.
ESRCH	R000104	The password is not correct or the user id is not completely defined (missing password or uid).
EACCES	R000104	The password is not correct or the user id is not completely defined (missing password or uid).
EINVAL	R000105	A bind argument is not valid.

For all of these reason codes, the LDAP return code will be LDAP_INVALID_CREDENTIALS.

If the SDBM database is to be used for authentication purposes only, consider having your clients use the authenticateOnly server control, to streamline bind processing. This supported control overrides any extended group membership searching and default group membership gathering and is supported for Version 3 clients. See Appendix G, "Supported server controls" on page 431 for more information.

Note that the SDBM backend only updates the default RACF on a given system. That is, the AT and **ONLYAT** clauses of the RACF commands, used to redirect RACF commands, are not exploited by SDBM.

See z/OS: Security Server RACF Command Language Reference for more information about the supported RACF commands.

See "Setting up for SDBM" on page 42 for information on getting your LDAP server configured with SDBM.

Note: The use of RACF passtickets is supported by the z/OS LDAP server. It is recommended that the LDAP server be run as a started task if RACF passticket support will be used. The job name associated with the LDAP Server started task should be used as the application name when generating RACF passtickets. Refer to z/OS: Security Server RACF Macros and Interfaces for more information about RACF passtickets.

Mapping LDAP-style names to RACF attributes

Following are tables that show the RACF attribute name and the corresponding LDAP-style attribute name for user (Table 30), group (Table 31 on page 209), and connection (Table 32 on page 210).

Table 30. Mapping of LDAP-style names to RACF attributes (user)

RACF segment name	RACF attribute name in altuser/adduser string	LDAP-style attribute name
User base or group base	DATA	racfInstallationData
User base or group base	MODEL	racfDatasetModel
User base	Not modifiable; listuser displays as CREATED	racfAuthorizationDate
User base	OWNER	racfOwner
User base	Multi-value: ADSP, SPECIAL, OPERATIONS, GRPACC, AUDITOR, OIDCARD, UAUDIT	racfAttributes
User base	PASSWORD	racfPassword
User base	Not modifiable - displayed as PASS-INTERVAL	racfPasswordInterval
User base	Not modifiable - displayed as racfPasswordChangeDate PASSDATE	
User base	NAME racfProgrammerName	
User base	DFLTGRP racfDefaultGroup	
User base	Not modifiable - displayed as racfLastAccess LAST-ACCESS	
User base	SECLEVEL	racfSecurityLevel
User base	ADDCATEGORY racfSecurityCategoryList	
User base	REVOKE racfRevokeDate	

Table 30. Mapping of LDAP-style names to RACF attributes (user) (continued)

RACF segment name	RACF attribute name in altuser/adduser string	LDAP-style attribute name
User base	RESUME racfResumeDate	
User base	WHEN(DAYS()) racfLogonDays	
User base	WHEN(TIME())	racfLogonTime
User base	CLAUTH	racfClassName
User base	GROUP	racfConnectGroupName
User base	AUTH not displayed by LDAP	racfConnectGroupAuthority
User base	UACC not displayed by LDAP	racfConnectGroupUACC
User base	SECLABEL	racfSecurityLabel
DFP segment - common to group or user	DATAAPPL	SAFDfpDataApplication
DFP segment - common to group or user	DATACLAS	SAFDfpDataClass
DFP segment - common to group or user	MGMTCLAS	SAFDfpManagementClass
DFP segment - common to group or user	STORCLAS	SAFDfpStorageClass
TSO segment	ACCTNUM	SAFAccountNumber
TSO segment	COMMAND	SAFDefaultCommand
TSO segment	DEST	SAFDestination
TSO segment	HOLDCLASS	SAFHoldClass
TSO segment	JOBCLASS	SAFJobClass
TSO segment	MSGCLASS	SAFMessageClass
TSO segment	PROC	SAFDefaultLoginProc
TSO segment	SIZE	SAFLogonSize
TSO segment	MAXSIZE	SAFMaximumRegionSize
TSO segment	SYSOUTCLASS	SAFDefaultSysoutClass
TSO segment	USERDATA	SAFUserdata
TSO segment	UNIT	SAFDefaultUnit
TSO segment	SECLABEL	SAFTsoSecurityLabel
LANGUAGE segment	PRIMARY	racfPrimaryLanguage
LANGUAGE segment	SECONDARY	racfSecondaryLanguage
CICS® segment	OPIDENT	racfOperatorIdentification
CICS segment	OPCLASS	racfOperatorClass
CICS segment	OPPRTY	racfOperatorPriority
CICS segment	XRFSOFF racfOperatorReSignon	
CICS segment	TIMEOUT	racfTerminalTimeout
OPERPARM segment	STORAGE	racfStorageKeyword
OPERPARM segment	AUTH	racfAuthKeyword
OPERPARM segment	MFORM	racfMformKeyword
OPERPARM segment	LEVEL	racfLevelKeyword

Table 30. Mapping of LDAP-style names to RACF attributes (user) (continued)

RACF segment name	RACF attribute name in altuser/adduser string	LDAP-style attribute name
OPERPARM segment	MONITOR	racfMonitorKeyword
OPERPARM segment	ROUTCODE	racfRoutcodeKeyword
OPERPARM segment	LOGCMDRESP	racfLogCommandResponseKeyword
OPERPARM segment	MIGID	racfMGIDKeyword
OPERPARM segment	DOM	racfDOMKeyword
OPERPARM segment	KEY	racfKEYKeyword
OPERPARM segment	CMDSYS	racfCMDSYSKeyword
OPERPARM segment	UD	racfUDKeyword
OPERPARM segment	MSCOPE	racfMscopeSystems
OPERPARM segment	ALTGRP	racfAltGroupKeyword
OPERPARM segment	AUTO	racfAutoKeyword
WORKATTR segment	WANAME	racfWorkAttrUserName
WORKATTR segment	WABLDG	racfBuilding
WORKATTR segment	WADEPT	racfDepartment
WORKATTR segment	WAROOM	racfRoom
WORKATTR segment	WAADDR1	racfAddressLine1
WORKATTR segment	WAADDR2	racfAddressLine2
WORKATTR segment	WAADDR3	racfAddressLine3
WORKATTR segment	WAADDR4	racfAddressLine4
WORKATTR segment	WAACCNT	racfWorkAttrAccountNumber
User OMVS segment	UID	racfOmvsUid
User OMVS segment	SHARED, AUTOUID	racfOmvsUidKeyword
User OMVS segment	HOME	racfOmvsHome
User OMVS segment	PROGRAM	racfOmvsInitialProgram
User OMVS segment	CPUTIMEMAX	racfOmvsMaximumCPUTime
User OMVS segment	ASSIZEMAX	racfOmvsMaximumAddressSpaceSize
User OMVS segment	FILEPROCMAX	racfOmvsMaximumFilesPerProcess
User OMVS segment	PROCUSERMAX	racfOmvsMaximumProcessesPerUID
User OMVS segment	THREADSMAX	racfOmvsMaximumThreadsPerProcess
User OMVS segment	MMAPAREAMAX	racfOmvsMaximumMemoryMapArea
Netview segment	IC	racfNetviewInitialCommand
Netview segment	CONSNAME	racfDefaultConsoleName
Netview segment	CTL	racfCTLKeyword
Netview segment	MSGRECVR	racfMessageReceiverKeyword
Netview segment	OPCLASS	racfNetviewOperatorClass
Netview segment	DOMAINS	racfDomains
Netview segment	NGMFADMN	racfNGMFADMKeyword
DCE segment	UUID	racfDCEUUID
DCE segment	DCENAME	racfDCEPrincipal

Table 30. Mapping of LDAP-style names to RACF attributes (user) (continued)

RACF segment name	RACF attribute name in altuser/adduser string	LDAP-style attribute name
DCE segment	HOMECELL	racfDCEHomeCell
DCE segment	HOMEUUID	racfDCEHomeCellUUID
DCE segment	AUTOLOGIN	racfDCEAutoLogin
User OVM segment	UID	racfOvmUid
User OVM segment	HOME	racfOvmHome
User OVM segment	PROGRAM	racfOvmInitialProgram
User OVM segment	FSROOT	racfOvmFileSystemRoot
LNOTES segment	SNAME	racfLNotesShortName
NDS segment	UNAME	racfNDSUserName
KERB segment	KERBNAME	krbPrincipalName
KERB segment	MAXTKTLFE	maxTicketAge
KERB segment	Not modifiable - displayed as KEY VERSION	racfCurKeyVersion
KERB segment	ENCRYPT	racfEncryptType
PROXY segment	BINDDN	racfLDAPBindDN
PROXY segment	BINDPW - value displayed is YES or NO	racfLDAPBindPw
PROXY segment	LDAPHOST	racfLDAPHost
EIM segment	LDAPPROF	racfLDAPProf

Table 31. Mapping of LDAP-style names to RACF attributes (group)

RACF segment name	RACF attribute name in altgroup/addgroup string	LDAP-style attribute name
User base or Group base	OWNER	racfOwner
User base or Group base	DATA	racfInstallationData
User base or Group base	MODEL	racfDatasetModel
Group base	SUPGROUP	racfSuperiorGroup
Group base	TERMUACC	racfGroupNoTermUAC
Group base	UNIVERSAL	racfGroupUniversal
Group base	Not modifiable - listgrp displays as SUBGROUP(S)	racfSubGroupName
Group base	Not modifiable - listgrp displays as USER(S)	racfGroupUserids
Group OMVS segment	GID	racfOmvsGroupId
Group OMVS segment	SHARED, AUTOGID	racfOmvsGroupIdKeyword
Group OVM segment	GID	racfOvmGroupId
DFP segment - common to group or user	DATAAPPL	SAFDfpDataApplication
DFP segment - common to group or user	DATACLAS	SAFDfpDataClass
DFP segment - common to group or user	MGMTCLAS	SAFDfpManagementClass
DFP segment - common to group or user	STORCLAS	SAFDfpStorageClass

Table 32. Mapping of LDAP-style names to RACF attributes (connection)

RACF segment name	RACF attribute name in connect string	LDAP-style attribute name
Connection base	Multi-value: ADSP, AUDITOR GRPACC, OPERATIONS, SPECIAL	racfConnectAttributes
Connection base	AUTHORITY	racfConnectGroupAuthority
Connection base	Not modifiable - displayed as CONNECT-DATE	racfConnectAuthDate
Connection base	Not modifiable - displayed as CONNECTS	racfConnectCount
Connection base	Not modifiable - displayed as LAST-CONNECT	racfConnectLastConnect
Connection base	OWNER	racfConnectOwner
Connection base	RESUME	racfConnectResumeDate
Connection base	REVOKE	racfConnectRevokeDate
Connection base	UACC	racfConnectGroupUACC

RACF namespace entries

When the SDBM database is used to make RACF information accessible over the LDAP protocol, the top four entries in the hierarchy are reserved, read-only, and generated by the server. The purpose of these reserved entries is to enable a hierarchical representation of RACF users, groups, and connections. Note that the root of the RACF namespace can be any DN and is not required to contain the sysplex attribute as the RDN. For example, the top four entries in Figure 34 are:

- cn=RACFA,o=IBM,c=US (suffixDN)
- profileType=User,cn=RACFA,o=IBM,c=US
- profileType=Group, cn=RACFA, o=IBM, c=US
- profileType=Connect,cn=RACFA,o=IBM,c=US

The value of the top DN is generated from the suffix line in the slapd.conf file for the SDBM database entry (see "Setting up for SDBM" on page 42).

Following is a high-level diagram of the RACF backend.

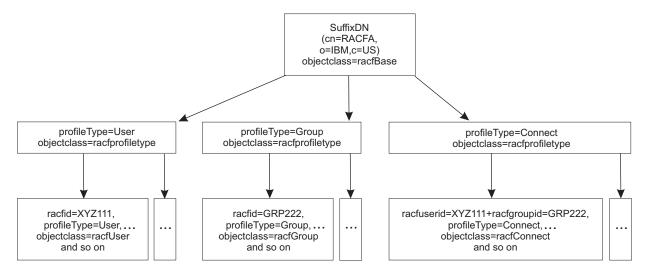


Figure 34. RACF namespace hierarchy

SDBM schema information

SDBM has an internal schema that it uses to check entries during add. The schema contains all the attributes and object classes that were in the slapd.at.racf and slapd.oc.racf schema files shipped in previous releases. The schema cannot be modified. The DN of the SDBM schema entry generated by the server is cn=schema, suffixDN. For example, if suffixDN is cn=RACFA, o=IBM, c=US then the schema entry DN would be

cn=schema,cn=RACFA,o=IBM,c=US

In this case, the SDBM schema can be published (displayed) by the following search command:

ldapsearch -h ldaphost -p ldapport -b "cn=SCHEMA,cn=RACFA,o=IBM,c=US" "objectclass=subschema"

See Appendix C, "SDBM schema" on page 397 for the complete SDBM schema.

When running TDBM and SDBM together, the TDBM schema will be used for initial DN normalization for selecting the appropriate backend. Since TDBM schema will be used for DN normalization, all schema elements used in formulating distinguished names for SDBM must appear in TDBM's schema definition.

SDBM support for pound sign

An SDBM DN can contain a pound sign (#) anywhere in the DN, including the suffix. The pound sign must be escaped by preceding it with a single backslash (\). Note that the suffix in the configuration file must use two backslashes (\\) to escape a pound sign, but only a single backslash is used in a DN.

For example, if the SDBM suffix in the configuration file is suffix sysplex=my\\#1plex

then the DN for the RACF user ab#c would be racfid=ab\#c,profiletype=user,sysplex=my\#1plex

Pound signs in DN values returned by SDBM from a search are always escaped by a single backslash.

When specifying a value containing a pound sign for an attribute within an add or modify request or within a search filter, do not escape the pound sign with a backslash, even if the value is a DN. For instance, to search for all RACF users starting with user#, use the search filter racfid=user#*. To add a user with default group d#1grp, specify:

racfdefaultgroup: racfid=d#1grp,profiletype=group,sysplex=my#1plex

within the entry.

SDBM operational behavior

Table 33 shows how the SDBM database behaves during different LDAP operations.

Table 33. RACF backend behavior

Target DN	LDAP operation behavior	
suffixDN	Add Error: Unwilling to perform	
	Modify Error: Unwilling to perform	
	Delete Error: Unwilling to perform	
	Modify DN Error: Unwilling to perform	
	Compare Compare attribute	
	Search base Return requested attributes	
	Search one level Perform a base search against each subordinate of this entry	
	Search subtree See "Searching the entire RACF database" on page 218	
	Bind Error: No credentials	
profiletype=User,suffixDN	Add Error: Unwilling to perform	
	Modify Error: Unwilling to perform	
	Delete Error: Unwilling to perform	
	Modify DN Error: Unwilling to perform	
	Compare Compare attribute	
	Search base Return requested attributes	
	Search one level See "Searching the entire RACF database" on page 218	
	Search subtree See "Searching the entire RACF database" on page 218	
	Bind Error: No credentials	

Table 33. RACF backend behavior (continued)

Target DN	LDAP operation behavior	
profiletype=Group,suffixDN	Add Error: Unwilling to perform	
	Modify Error: Unwilling to perform	
	Delete Error: Unwilling to perform	
	Modify DN Error: Unwilling to perform	
	Compare Compare attribute	
	Search base Return requested attributes	
	Search one level See "Searching the entire RACF database" on page 218	
	Search subtree See "Searching the entire RACF database" on page 218	
	Bind Error: No credentials	
profiletype=Connect,suffixDN	Add Error: Unwilling to perform	
	Modify Error: Unwilling to perform	
	Delete Error: Unwilling to perform	
	Modify DN Error: Unwilling to perform	
	Compare Compare attribute	
	Search base Return requested attributes	
	Search one level See "Searching the entire RACF database" on page 218	
	Search subtree See "Searching the entire RACF database" on page 218	
	Bind Error: No credentials	

Table 33. RACF backend behavior (continued)

Target DN	LDAP operation behavior	
racfid=XYZ111,profiletype=User, suffixDN	Add	Perform an adduser RACF command using USER=XYZ111
SullixBiv	Modify	Perform an altuser RACF command using USER=XYZ111
	Delete	Perform a deluser RACF command using USER= XYZ111
	Modify DN Error: Unwilling to perform	
	Compa	re Compare requested attribute with data returned from listuser RACF command using USER=XYZ111
	Search	base Perform a listuser RACF command using USER=XYZ111
	Search	one level Empty search results (this is a leaf node in the hierarchy)
	Search	subtree Perform a listuser RACF command using USER=XYZ111
	Bind	If bind type is not simple, error: Unwilling to perform, else usepasswd() to verify the user ID and password combination and perform a listuser RACF command using USER=XYZ111
racfid=GRP222,profiletype= Group, suffixDN	Add	Perform an addgroup RACF command using GROUP=GRP222
0.00ap, 0.0	Modify	Perform an altgroup RACF command using GROUP=GRP222
	Delete	Perform a delgroup RACF command using GROUP=GRP222
	Modify DN Error: Unwilling to perform	
	Compa	re Compare requested attribute with data returned from listgrp RACF command using GROUP=GRP222
	Search	base Perform a listgrp RACF command using GROUP=GRP222
	Search	one level Empty search results (this is a leaf node in the hierarchy)
	Search	subtree Perform a listgrp RACF command using GROUP=GRP222
	Bind	Error: No credentials

Table 33. RACF backend behavior (continued)

Target DN	LDAP o	peration behavior
racfuserid=XYZ111+racfgroupid= GRP222,profiletype=Connect, suffixDN	Add	Perform a connect RACF command for USER=XYZ111 using GROUP=GRP222
	Modify	Perform a connect RACF command for USER=XYZ111 using GROUP=GRP222
	Delete	Perform a remove RACF command for USER=XYZ111 using GROUP=GRP222
	Modify DN Error: Unwilling to perform	
	Compare Compare requested attribute with data returned from listuser RACF command using USER=XYZ111 Search base Perform a listuser RACF command using USER=XYZ111 Search one level Empty search results (this is a leaf node in the hierarchy)	
	Search	subtree Perform a listuser RACF command using USER=XYZ111
	Bind	Error: No credentials

If LDAP is running with an SDBM backend, the Idap modify and Idap add APIs can return LDAP OTHER or LDAP SUCCESS and have completed a partial update to an entry in RACF. The results will match what would occur if the update were done using the RACF altuser, altgroup, and connect commands. If several RACF attributes are being updated and one of them is in error, RACF reports on the error, but still updates the other attributes. The RACF message text is also returned in the result.

The RACF connect command is used to both add a user connection to a group and to modify a user's connection to a group. As a result, the SDBM add and modify support for connection entries is different than normal LDAP support:

- · When adding a connection entry that already exists, the entry is modified using the specified attributes. There is no indication returned that the entry already existed.
- When modifying a connection entry that does not exist, the entry is added using the specified attributes. There is no indication returned that the entry did not exist.

Notes about specifying attribute values:

1. In LDAP, the format of the value of the Kerberos principal name attribute, krbPrincipalName, is userid@realm. In SDBM, the userid portion of the name is case-sensitive while the realm portion of the name is not. In addition, SDBM only operates on the RACF local realm. If the realm specified in the value is not the local realm, the operation fails. For example, when searching in SDBM for a user entry using the search filter krbPrincipalName=krbuser1@myrealm.com, the search fails if myrealm.com is not the RACF local realm.

To facilitate using the **krbPrincipalName** attribute, the attribute value can be specified without the @realm portion. In this case, the realm is assumed to be the RACF local realm. For example, when adding a user entry with krbuser1 as the userid portion of that Kerberos principal name, the krbPrincipalName attribute can be specified as

krbPrincipalName=krbuser1

krbPrincipalName=krbuser1@myrealm.com

where myrealm.com is the RACF local realm.

The **krbPrincipalName** value returned by SDBM from a search is always the complete principal name, userid@realm, where realm is the RACF local realm.

2. There are several SDBM attributes whose value is a RACF user or group name. For convenience, this value can be specified either as just the RACF name or as the complete LDAP DN. For example, when adding a user with a default group of grp222, the racfDefaultGroup attribute can be specified as racfDefaultGroup=grp222

racfDefaultGroup=racfid=grp222,profiletype=group,sysplex=myplex

where sysplex=myplex is the SDBM suffix.

The value returned by SDBM from a search is always the complete LDAP DN.

- 3. For multi-value attributes, the RACF altuser command does not always support the ability to both add a value and replace the existing value. As a result, SDBM does not always respect the type of modification (add versus replace) that is specified in a modify command. Values for the following multi-value attributes are always added to the existing value (even if replace is specified): racfAttributes, racfClassName, racfConnectAttributes, racfLevelKeyword, racfMformKeyword, racfMonitorKeyword, racfSecurityCategoryList. Values for the following multi-value attributes always replace the existing value (even if add is specified); racfDomains, racfMscopeSystems. racfNetviewOperatorClass, racfOperatorClass, racfRoutcodeKeyword. Values for the following multi-value attributes either are added to the existing values or replace the existing values, depending on the new and existing values: racfAuthKeyword.
- 4. In order to update CICS-related attributes, CICS must be set up on your system; otherwise, errors result.
- 5. For modify, if a request is made to delete a specific attribute value for an attribute where specific values cannot be selectively deleted, LDAP_UNWILLING_TO_PERFORM is returned. There are three attributes where specific attribute values are accepted: racfAttributes, racfSecurityCategoryList, and racfConnectAttributes. If an attempt is made to delete any attribute that has no corresponding delete command in RACF, LDAP_UNWILLING_TO_PERFORM is returned.

Mapping SDBM requests to RACF commands

SDBM uses the RACF R_Admin interface to issue RACF commands to process LDAP requests. The RACF commands are issued under the identity of the bind user. See z/OS: Security Server RACF Command Language Reference for more information on the authorization required to use each RACF command. Following is a mapping of the RACF commands issued by SDBM for each type of LDAP request.

- 1. Add, modify, delete, compare, and bind an entry for user U result in adduser u, altuser u, deluser u, listuser u, and listuser u respectively.
- 2. Add, modify, delete, and compare an entry for group G result in addgroup g, altgroup g, delgroup g, and **listgrp** g respectively.
- 3. Add, modify, delete, compare an entry for the connection of user U to group G result in connect u group(g), connect u group(g), remove u group(g), and listuser u respectively.
- 4. Search:
 - a. Searching for a user entry results in either listuser value, search class(user) filter(value), getpwuid(value), or search class(user) uid(value)
 - b. Searching for a group entry results in either listgrp value, search class(group) filter(value), getgrgid(value), or search class(user) gid(value)

c. Searching for a connect entry results in listuser value, or listgrp value, or search class(group) filter(value) followed by listgrp group for each group returned by search

where value is the value specified in the search filter.

SDBM search capabilities

SDBM supports a limited set of search filters. Following is a list of the filters and the type of entry that matches each filter:

Table 34. SDBM-supported search filters

Filter	Matching entries
objectclass=*	All entries for RACF users, groups, and connections
racfid= <any_value></any_value>	User or group entries for the RACF users and groups with the specified name (can contain wildcards)
racflnotesshortname= <any_value></any_value>	User entry for the RACF user with the specified LNOTES SNAME
racfndsusername= <any_value></any_value>	User entry for the RACF user with the specified NDS UNAME
racfomvsuid= <number></number>	User entry for a RACF user with the specified OMVS UID
racfomvsuid;allOMVSids= <number></number>	User entries for all the RACF users with the specified OMVS UID
racfomvsgroupid= <number></number>	Group entry for a RACF group with the specified OMVS GID
racfomvsgoupid;allOMVSids= <number></number>	Group entries for all the RACF groups with the specified OMVS GID
krbprincipalname= <any_name></any_name>	User entry for the RACF user with the specified KERB KERBNAME
racfuserid= <any_value></any_value>	Connection entries for the RACF users with the specified name (can contain wildcards)
racfgroupid= <any_value></any_value>	Connection entries to the RACF groups with the specified name (can contain wildcards)
(&(racfuserid= <any_value>)(racfgroupid= <any_value>))</any_value></any_value>	Connection entries for the RACF users with the first specified name to the RACF groups with the second specified name (both can contain wildcards)

Table 35 shows the search filters that are supported for each search base.

Table 35. RACF backend search filters

Search base	Filters supported
suffixDN	objectclass=*
(root of the directory)	racfid= <any_value></any_value>
	racfInotesshortname= <any_value></any_value>
	racfndsusername= <any_value></any_value>
	racfomvsuid=< <i>number</i> >
	racfomvsuid;allOMVSids= <number></number>
	racfomvsgroupid= <number></number>
	racfomvsgroupid;allOMVSids= <number></number>
	krbprincipalname= <any_name></any_name>
	racfuserid= <any_value></any_value>
	racfgroupid= <any_value></any_value>
	(&(racfuserid= <any_value>)(racfgroupid=</any_value>
	<any_value>))</any_value>

Table 35. RACF backend search filters (continued)

Search base	Filters supported
profileType=user,suffixDN	objectclass=* racfid= <any_value> racfinotesshortname=<any_value> racfndsusername=<any_value> racfomvsuid=<number> racfomvsuid;allOMVSids=<number> krbprincipalname=<any_value></any_value></number></number></any_value></any_value></any_value>
profileType=group,suffixDN	objectclass=* racfid=< <i>any_value</i> > racfomvsgroupid=< <i>number</i> > racfomvsgroupid;allOMVSids=< <i>number</i> >
profileType=connect,suffixDN	objectclass=* racfuserid=< <i>any_value</i> > racfgroupid=< <i>any_value</i> > (&(racfuserid=< <i>any_value</i> >)(racfgroupid= < <i>any_value</i> >))
racfid=abc,profileType=user,suffixDN	objectclass=*
racfid=xyz,profileType=group,suffixDN	objectclass=*
racfuserid=abc+racfgroupid=xyz,profileType=connect, suffixDN	objectclass=*

Except for the AND filter for connections, complex search filters that include NOT, AND, OR, LE, or GE constructs are not supported.

The values for the racfid, racfuserid, and racfgroupid filters can include the wild cards supported by RACF. These wild cards are '*' which represents any number of characters, and '%' which represents one character. For example:

(&(racfuserid=usr*)(racfgroupid=*grp))

searches for all the connections between users whose names begin with usr and groups whose names end with grp.

Note about searching universal groups: Most of the members of a RACF universal group are not actually contained in the group's list of members. As a result, a search of the entry for a universal group does not return most of the group's members. In addition, a search for the connection entry corresponding to a member of a universal group can return different results depending on the connection search filter that is used:

- If the racfuserid part of the connection search filter does not contain a wild card, then the connection entry is returned for the specified racfuserid.
- If the racfuserid part of the connection search filter contains a wild card, then the connection entry for a user is returned only if the user is explicitly contained in the universal group's list of members.

Searching the entire RACF database

Most searches that query the entire RACF database, for example, a subtree search from one of the top four directory entries, return only the DN (distinguished name) attribute. You may then obtain more specific data about a particular user or group on a follow-up search using a specific DN as the search base. See "RACF-style distinguished names" on page 156 for more information.

The exceptions to this are searches using the "application ID" filters:

racflnotesshortname=<anv value> racfndsusername=<any value> krbprincipalname=<any name> racfomvsuid=<number> racfomvsgroupid=<number>

Because these searches can match only a single RACF user, the entire user entry is returned in the search results.

RACF restriction on amount of output: When processing LDAP bind, compare, and search requests, SDBM uses the RACF R Admin interface to issue RACF search, listuser, and listgrp commands. R Admin limits the number of records in its output to 4096. This means that the RACF search command output may be incomplete if you have a large number of users, groups, and connections. Similarly, the RACF listuser and listgrp output may be incomplete for a user who is a member of many groups and for a group which has many members. If this occurs, the entries returned by SDBM will also be incomplete. In particular, when binding with a such a user, the list of groups associated with the user will be incomplete.

Using SDBM to change a user password in RACF

There are two ways to use SDBM to change a user password in RACF.

1. The user password of the bind user can be changed during an LDAP simple bind to SDBM. The simple bind can occur as part of an LDAP function such as search, add, or modify. The password change is provided in the password portion of the LDAP simple bind. The password must be in the following format:

password/newpassword

The forward slash ('/') is used as the indication of a password change during the LDAP simple bind. Password changes made using the LDAP simple bind to the SDBM backend of the z/OS LDAP server are subject to the system password rules. A password change will fail with LDAP return code LDAP INVALID CREDENTIALS and LDAP reason code of:

```
R000101 The new password is not valid.
```

if the new password does not pass the rules established on the system.

Note that once the bind succeeds, the password is changed even if the LDAP function eventually fails.

For example, the following command changes the password for RACF user U1 from abc to xyz, assuming the SDBM suffix is sysplex=sysplexa:

```
ldapsearch -h ldaphost -p ldapport -D racfid=u1,profiletype=user,sysplex=sysplexa
  -w abc/xyz -s base -V 3 -b "" objectclass=*
```

2. To change any user's password, create an LDIF file that modifies the racfPassword attribute for that user and then invoke **Idapmodify** to change the password. If the syntax of the new password is not valid, the command fails, returning "ldap_modify: Unknown error". (Note that this response can also be returned under other circumstances.)

For example, the following LDIF file, pw.mod, resets the password for RACF user U1 to xyz, assuming the SDBM suffix is sysplex=sysplexa. The racfAttributes: noexpired record is added to result in a new password that is not expired. If noexpired is not specified, then the password is reset but is expired, requiring U1 to change the password at next logon.

```
dn: racfid=u1,profiletype=USER,sysplex=sysplexa
changetype: modify
add: x
racfpassword: xvz
racfattributes: noexpired
```

Then, assuming that the RACF user admin1 has the necessary RACF authorization to update RACF, the command:

```
ldapmodify -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa
  -w passwd -f pw.mod
```

modifies the password for U1.

Using LDAP operation utilities with SDBM

The LDAP operation utilities described in the z/OS: Security Server LDAP Client Programming can be used to update data in RACF. Following are some examples. These examples assume that the RACF user admin1 has the necessary RACF authorization to make these RACF updates and that sysplex=sysplexa is the SDBM suffix.

Example: displaying the SDBM schema

- To see the contents of the schema used by SDBM, the following search command can be performed:
- ldapsearch -h ldaphost -p ldapport -b "cn=schema,sysplex=sysplexa" "objectclass=subschema"
- The beginning of the results of the search will look like:

```
CN=SCHEMA, sysplex=sysplexa
cn=SCHEMA
subtreespecification=NULL
attributetypes=( 2.5.21.5 NAME 'attributeTypes' EQUALITY objectIdentifierFirstComponentMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.3 USAGE directoryOperation )
```

See Appendix C, "SDBM schema" on page 397 for the complete SDBM schema.

Example: adding a user to RACF

If the LDIF file user.add contains:

```
dn: racfid=newuser,profiletype=user,sysplex=sysplexa
objectclass: racfUser
racfid: newuser
```

The following command will add user ID newuser to RACF:

```
ldapadd -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w
passwd -f user.add
```

Note that the only required attribute to add a user is the user ID specified as racfid. This mimics the RACF adduser command.

Example: modifying a user in RACF

To add a TSO segment for newuser, the LDIF file user.mods could contain:

```
dn: racfid=newuser,profiletype=user,sysplex=sysplexa
changetype: modify
objectclass: SAFTsoSegment
SAFAccountNumber: 123
SAFHoldClass: H
SAFLogonSize: 1024
```

The command:

```
ldapmodify -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w
passwd -f user.mods
```

modifies the RACF user profile for user ID newuser, adding a TSO segment with the specified values.

Example: searching for user information in RACF

To see the information in RACF for newser, the following search command can be performed:

```
ldapsearch -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w passwd
 -b "racfid=newuser,profiletype=user,sysplex=sysplexa" "objectclass=*
```

The results that are returned are most of the data that RACF displays on a **listuser** command, but using LDAP-style attribute names. Following is an example for newuser:

```
racfid=NEWUSER,profiletype=USER,sysplex=sysplexa
objectclass=racfUser
objectclass=racfBaseCommon
objectclass=SAFTsoSegment
racfid=NEWUSER
racfprogrammername=UNKNOWN
racfowner=racfid=ADMIN1,profiletype=USER,sysplex=sysplexa
racfauthorizationdate=98.001
racfdefaultgroup=racfid=GROUPID,profiletype=GROUP,sysplex=sysplexa
racfpasswordchangedata=00.000
racfpasswordinterval=90
racfattributes=NONE
racfrevokedata=NONE
racfresumedate=NONE
racflastaccess=UNKNOWN
racfclassname=NONE
racfinstallationdata=NO-INSTALLATION-DATA
\verb|racfdataset| model = \verb|NO-MODEL-NAME| \\
racflogondays=ANYDAY
racflogontime=ANYTIME
racfconnectgroupname=racfid=GROUPID,profiletype=GROUP,sysplex=sysplexa
racfsecuritylevel=NONE SPECIFIED
racfsecuritycategorylist=NONE SPECIFIED
racfsecuritylabel=NONE SPECIFIED
safaccountnumber=123
safholdclass=H
saflogonsize=00001024
safmaximumregionsize=00000000
safuserdata=0000
```

Example: adding a group to RACF

If the LDIF file group.add contains:

dn: racfid=grp222,profiletype=group,sysplex=sysplexa objectclass: racfGroup racfid: grp222

The following command adds group ID grp222 to RACF:

```
ldapadd -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w passwd
  -f group.add
```

Note that the only required attribute to add a group is the group ID specified as racfid. This mimics the RACF addgroup command.

The commands for modifying, searching, and removing a RACF group using SDBM are very similar to the corresponding commands for a RACF user. See the examples in this section for a RACF user for more information.

Example: connecting a user to a group in RACF

To connect newuser to group grp222, the LDIF file connect.add could contain:

dn: racfuserid=newuser+racfgroupid=grp222,profiletype=connect,sysplex=sysplexa objectclass: racfconnect racfuserid: newuser racfgroupid: grp222

The command:

```
ldapadd -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w passwd
  -f connect.add
```

makes newuser a member of the grp222 group. Note that grp222 must be an existing RACF group ID, newuser must be an existing RACF user ID, and the only required attributes to add a connection are racfuserid (the user ID) and racfgroupid (the group ID).

Example: searching for information about a user's connection to a group in RACF To see information about newser's connection to the grp222 group, the following search can be performed:

```
ldapsearch -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w passwd
  -b "racfuserid=newuser+racfgroupid=grp222,profiletype=connect,sysplex=sysplexa" "objectclass=*"
```

The result returned is the information from the GROUP section that RACF displays on a listuser command, but using LDAP-style attribute names. Following is an example for newser's connection to grp222:

racfuserid=NEWUSER+racfgroupid=GRP222,profiletype=CONNECT,sysplex=sysplexa objectclass=racfConnect racfuserid=NEWUSER racfgroupid=GRP222 racfconnectgroupauthority=USE racfconnectowner=racfid=ADMIN1,profile=USER,sysplex=sysplexa racfconnectauthdate=00.224 racfconnectcount=00 racfconnectgroupuacc=NONE racfconnectlastconnect=UNKNOWN racfconnectattributes=NONE racfconnectrevokedate=NONE racfconnectresumedate=NONE

To see all the groups that newuser is connected to, either of the following searches can be performed:

```
ldapsearch -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w passwd
  -b "profiletype=connect,sysplex=sysplexa" "racfuserid=newuser"
```

```
ldapsearch -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa -w passwd
  -b "profiletype=connect,sysplex=sysplexa" "(&(racfuserid=newuser)(racfgroupid=*))"
```

For both commands, the results are:

```
racfuserid=NEWUSER+racfgroupid=GROUPID,profiletype=CONNECT,sysplex=sysplexa
```

```
racfuserid=NEWUSER+racfgroupid=GRP222,profiletype=CONNECT,sysplex=sysplexa
```

Note that GROUPID was the default group to which newuser was connected when newuser was created.

Example: removing a user from a group in RACF

The following command removes newser from the grp222 group (the equivalent of the RACF remove command):

```
ldapdelete -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa
  -w passwd "racfuserid=newuser+racfgroupid=grp222,profiletype=connect,sysplex=sysplexa"
```

Example: removing a user from RACF

The following command removes the newuser user profile from RACF, also removing all of newuser's connections to groups (the equivalent of a RACF deluser command):

```
ldapdelete -h ldaphost -p ldapport -D racfid=admin1,profiletype=user,sysplex=sysplexa
  -w passwd "racfid=newuser,profiletype=user,sysplex=sysplexa"
```

Deleting attributes

The racfAttributes and racfConnectAttributes attributes are treated as multi-valued attributes in LDAP and represent the flags in the USER BASE RACF segment and the CONNECT BASE RACF segment, respectively. The available values for racfAttributes are:

- ADSP
- AUDITOR
- GRPACC
- OIDCARD
- OPERATIONS
- SPECIAL
- UAUDIT

The available values for racfConnectAttributes are:

- ADSP
- AUDITOR
- GRPACC
- OPERATIONS
- SPECIAL

If a request is made to delete the racfAttributes or racfConnectAttributes attribute and no values are provided, SDBM generates commands to delete each possible racfAttributes or racfConnectAttributes attribute value (even if the user or connection does not currently have that value). Deleting a specific value for racfAttributes or racfConnectAttributes requires that the value itself be specified on the delete operation. For example, to remove the OPERATIONS and AUDITOR values from the racfAttributes values of user ID user1 (leaving any other racfAttributes values the user has), you would issue an **Idapmodify** command with the following file:

dn: racfid=user1,profiletype=user,sysplex=sysplexa changetype: modify

delete: racfAttributes racfAttributes: OPERATIONS racfAttributes: AUDITOR

To remove all the racfAttributes values of user ID user1, you would issue an **Idapmodify** command with the following file:

dn: racfid=user1,profiletype=user,sysplex=sysplexa changetype: modify

delete: racfAttributes

Following are some additional examples of deleting attributes:

• dn: racfid=user1,profiletype=user,sysplex=sysplexa changetype: modify

delete: racfProgrammerName

Returns: LDAP_UNWILLING_TO_PERFORM

The **racfProgrammerName** attribute is one that cannot be deleted.

• dn: racfid=user1,profiletype=user,sysplex=sysplexa

changetype: modify delete: racfBuilding racfBuilding: 001

Returns: LDAP UNWILLING TO PERFORM

You cannot specify a value to be removed for racfBuilding.

• dn: racfid=user1,profiletype=user,sysplex=sysplexa changetype: modify delete: racfBuilding

Expected result: successful removal of the attribute racfBuilding and LDAP_SUCCESS returned.

Chapter 19. Kerberos authentication

The z/OS LDAP server allows clients to authenticate to the server by using IBM's Network Authentication and Privacy Service which is better known as Kerberos Version 5. Kerberos is a trusted third party, private-key, network authentication system. In Kerberos, a ticket, a packet of information used by a client to prove its identity, is passed to a server in place of a user name and password. This ticket is encrypted and cannot be duplicated. After the server verifies the client ticket, it sends its own ticket to the client in order for the client to authenticate it. Once the mutual authentication process is complete, the client and server have authenticated each other.

In the z/OS LDAP server, Kerberos is used for authentication only. The Kerberos options for integrity and confidentiality are not supported. Authorization information for ACLs is gathered by the LDAP server after the authentication process has completed and is based on the Kerberos identity of the bound client. For more information about Kerberos refer to z/OS: Security Server Network Authentication Service Administration and z/OS: Security Server Network Authentication Service Programming .

Setting up for Kerberos

Kerberos Version 5 binds, defined in IETF RFC 2222, are performed using the Generic Security Services Application Programming Interface (GSS API) defined in IETF RFCs 2743 and 2744. From this point on the phrase "GSS API bind" is used to refer to Kerberos Version 5 binds. Before attempting to perform a Kerberos GSS API bind, be sure to:

- 1. Have the Network Authentication and Privacy Service (Kerberos 5) installed and configured and the service started.
- 2. Create a Kerberos identity for the user ID that will start the LDAP server. For example: ALTUSER LDAPSRV PASSWORD(password) NOEXPIRED KERB(KERBNAME(LDAP/hostname))
 - Note that the word "LDAP" must be uppercase in the **kerbname** section of the **altuser** command. Also, the *hostname* needs to be the primary *hostname* for the system in DNS.
- 3. If the KDC (Key Distribution Center) is not located on the same machine as the LDAP server, you have to generate a keytab file for the server. To generate a keytab for the server, issue the following commands:
 - a. First check the version of the server's Kerberos key (this is necessary since the version is updated every time the password is changed):
 - LISTUSER LDAPSRV NORACF KERB
 - b. Now the **keytab** command can be issued from the z/OS shell with the version from the **LISTUSER** command:

```
keytab add LDAP/hostname -p password -v 001
```

The **-k** *filename* option may also be used if you want to use your own keytab file rather than the Kerberos default keytab file. It is also important to note that when issuing Kerberos commands all passwords must be in uppercase.

If the KDC and LDAP server are on the same system, you do not need a keytab file. If the ID which starts the LDAP server has READ access to the IRR.RUSERMAP facility class in RACF, then this can be used instead of a keytab file. Following are the RACF commands to do this:

```
RDEFINE FACILITY IRR.RUSERMAP UACC(NONE)
PERMIT IRR.RUSERMAP CLASS(FACILITY) ID(LDAPSRV) ACCESS(READ)
SETR RACLIST(FACILITY) REFRESH
```

4. Enable your configuration file for Kerberos authentication.

```
# Global Section
supportKrb5 yes
serverKrbPrinc LDAP/hostname@MYREALM.COM
```

krbLDAPAdmin ibm-kn=ldapadm@MYREALM.COM krbKeytab none # TDBM Section krbIdentityMap on # SDBM Section krbIdentityMap on

Note: The "LDAP" portion of the serverKrbPrinc identity must be uppercase in the configuration file and in the Kerberos segment of the RACF ID where it is defined. Check your KDC for case requirements.

5. Start your server. Your LDAP server is now configured with Kerberos support.

Updating the directory schema for Kerberos

In order to enable Kerberos GSS API Authentication, the directory schema must contain the Kerberos related schema elements which are defined in schema.user.ldif. For more information on schema updates see "Updating the schema" on page 173.

Table 36 lists the Kerberos object classes and attributes.

Table 36. Kerberos attributes and object classes

Attribute	Object Class	Description
krbRealmName-V2	krbRealm-V2	This attribute represents the Kerberos Realms of which entries in the LDAP server are members. The entry that contains this attribute also contains the krbPrincSubtree attribute.
krbPrincSubtree	krbRealm-V2	This attribute is in the same entry as the krbRealmName-V2 attribute and it identifies the directory subtrees where entries may contain Kerberos information.
krbPrincipalName	(no object class)	The attribute is used to define the entry's Kerberos identity. This attribute is used for identity mapping. Currently this attribute is not associated with an object class. This means that for an entry to contain this attribute you can add the object class extensibleObject or define and add your own object class.
krbAliasedObjectName	krbAlias	This attribute allows an entry to be mapped to another entry's DN.
krbHintAliases	krbAlias	This attribute is used as an authorization list. If another entry's DN is in this list and that entry specified this entry as a krbAliasedObjectName then the mapping is allowed.
altSecurityIdentities	ibm-securityIdentities	If a user is defined to a case-insensitive Kerberos server, then the Kerberos identity associated with this entry is stored as an altSecurityIdentity rather than a krbPrincipalName.
ibm-kn	(no object class)	This attribute is a pseudo-DN so that Kerberos identities can be represented as DNs for access control. Currently this attribute is not associated with an object class. This means that for an entry to contain this attribute you can add the object class extensibleObject or define and add your own object class.

Identity mapping

The following sections describe the mapping that is done depending on your configuration. After all the identity mapping takes place you are left with a list of DNs that are used for access control and group gathering.

Default mapping

The GSS API bind operation passes a Kerberos identity to the LDAP server which in its initial form cannot be used for access control in the server. This Kerberos identity known as cprincipal0REALMis converted to a DN of the form ibm-kn=<principal>@<REALM>. Now this Kerberos DN can be used in access control lists.

For example, if you performed a Kerberos bind as jeff@IBM.COM you would be mapped to ibm-kn=jeff@IBM.COM and this DN is added to a list of DNs that will be used for access control throughout the server. This is known as the default mapping and is always performed when a SASL bind with a mechanism of GSS API is performed.

TDBM mapping

Another form of mapping is to map the Kerberos identity to TDBM DNs. The following algorithm is used to perform this type of identity mapping if the krbldentityMap configuration option is on for this backend.

- 1. Search the entire TDBM backend for the realm entry that corresponds to the Kerberos identity. This is done by searching for objectclass = krbRealm and krbRealmName-V2 = <REALM>, where <REALM> is the realm portion of the bound Kerberos identity. If the realm is found in the directory, then all of its krbPrincSubtree values are gathered for use in the next part of this algorithm.
- 2. If krbPrincSubtree values exist, then each subtree will be searched for the entry or entries that contain the attribute

```
krbPrincipalName = <pri>principal>@<REALM>
```

where <principal>@<REALM> is the bound Kerberos identity.

- 3. If an entry or entries were found in step 2 with the correct krbPrincipalName, their DNs are added to the DN list. If the krbAliasedObjectName attribute exists in the entry that was found, then more work needs to be done. The entry specified as a krbAliasedObjectName must allow this entry to use its DN. So the entry specified in the krbAliasedObjectName must have the DN of the entry in its list of krbHintAliases. If it does, then the krbAliasedObjectName value is added to the DN list.
- 4. The final step is to search the entire database for entries that have an object class

```
objectclass = ibm-securityIdentities
```

```
and the attribute
```

altSecurityIdentities = KERBEROS:<principal>@<REALM>

where <principal>@<REALM> is the bound Kerberos identity.

SDBM mapping

If an SDBM backend is configured and the krbldentityMap configuration is ON, then the SDBM backend tries to map the Kerberos identity to the appropriate RACF ID. If a RACF ID is found, then the SDBM DN that represents the RACF ID is added to the list of DNs.

Configuring access control

Since we now have a list of alternate DNs, access control has been changed to operate on the list of DNs rather than just a single DN. Group gathering is also performed on all of the DNs in the list. The following examples show how access control could be configured for Kerberos binds.

1. Setting up new ACLs in your directory:

Use ibm-kn=<principal>@<REALM> for your aclEntry values.

Example:

```
dn: cn=Scott,o=IBM,c=US
aclEntry: access-id:ibm-kn=jeff@IBM.COM:normal:r
```

If jeff@IBM.COM performed a Kerberos bind to the server, he will be mapped to ibm-kn=jeff@IBM.COM and he would get read access to normal data in the Scott entry.

- 2. Use existing ACLs (Method 1). This method is used for Kerberos identities that are defined to IBM KDCs or case-sensitive KDCs.
 - a. Set up and add the realm entry in the database.

Example:

dn: krbRealmName-V2=IBM.COM,o=IBM,c=US objectclass: krbRealm krbRealmName-V2: IBM.COM krbPrincSubtree: o=IBM,c=US

This example states that if a bound Kerberos identity has a realm of IBM.COM, then identity mapping is performed in the o=IBM, c=US subtree.

b. Add the **krbPrincipalName** attribute to your entries.

Example:

```
dn: cn=Jeff,o=IBM,c=US
objectclass: extensibleObject
krbPrincipalName: jeff@IBM.COM
```

In this example, the realm object for jeff@IBM.COM is found and the o=IBM,c=US subtree is searched for krbPrincipalName = jeff@IBM.COM. Because there is no krbAliasedObjectName attribute in the Jeff entry, only the DN cn=Jeff,o=IBM,c=US is added to the DN list along with the default mapping of ibm-kn=jeff@IBM.COM.

Therefore, if cn=Jeff,o=IBM,c=US was already defined in another entry's aclEntry, then jeff@IBM.COM will still have that access to the entry. For example:

```
dn: cn=Ken,o=IBM,c=US
aclEntry: access-id:cn=Jeff,o=IBM,c=US:normal:w
```

In this example jeff@IBM.COM will still maintain access to the cn=Ken,o=IBM,c=US entry since TDBM mapping was performed.

c. The **krbAliasedObjectName** can also be used for identity mapping.

Example:

```
dn: cn=Jeff,o=IBM,c=US
objectclass: extensibleObject
objectClass: krbAlias
krbPrincipalName: jeff@IBM.COM
krbAliasedObjectName: cn=Tim,o=IBM,c=US
```

In this example, the realm object for jeff@IBM.COM is found and the o=IBM,c=US subtree is searched for krbPrincipalName = jeff@IBM.COM. The search results in cn=Jeff.o=IBM.c=US being added to the DN list. Because there is a krbAliasedObjectName attribute in the Jeff entry, we need to look at the Tim entry before we add cn=Tim,o=IBM,c=US to the DN list. In order to use Tim's DN for access control he must authorize Jeff to do so. Tim's entry must look like the following:

```
dn: cn=Tim.o=IBM.c=US
objectclass: krbAlias
krbHintAliases: cn=Jeff,o=IBM,c=US
```

Since Tim has Jeff listed as a krbHintAliases, the value of krbAliasedObjectName cn=Tim,o=IBM,c=Us can be added to the DN list. If the Tim entry did not contain the krbHintAliases with Jeff as its value, then Tim's DN would not be added to the DN list.

Therefore, if cn=Tim,o=IBM,c=US was already defined in another entry's aclEntry then jeff@IBM.COM will still have that access to the entry. For example:

```
dn: cn=Kim.o=IBM.c=US
aclEntry: access-id:cn=Tim,o=IBM,c=US:normal:w
```

In this example, jeff@IBM.COM still maintains write access to the Kim entry since TDBM mapping was performed and Jeff was aliased to Tim.

3. Use existing ACLs (Method 2). This method should be used for case-insensitive KDCs. Set up your TDBM entries with the altSecurityIdentities attribute.

Example:

```
dn: cn=Jeff,o=IBM,c=US
objectclass: ibm-securityIdentities
altSecurityIdentity: KERBEROS:jeff@IBM.COM
```

Now if jeff@IBM.COM performs a Kerberos bind he will be mapped to ibm-kn=jeff@IBM.COM as well as cn=Jeff,o=IBM,c=US.

4. Therefore, if cn=Jeff,o=IBM,c=US was already defined in another entry's aclEntry then jeff@IBM.COM still has that access to the entry.

For example:

```
dn: cn=Ken,o=IBM,c=US
aclEntry: access-id:cn=Jeff,o=IBM,c=US:normal:w
```

In this example jeff@IBM.COM still maintains write access to the Ken entry since TDBM mapping was performed.

Example of setting up a Kerberos directory

The following diagram shows an example of how you could set up a Kerberos directory.

Note: Due to space limitations of the diagram, the entries in the example do not contain all of the necessary information to make them valid directory entries. For example, object classes and required attributes have been left out of many of the entries.

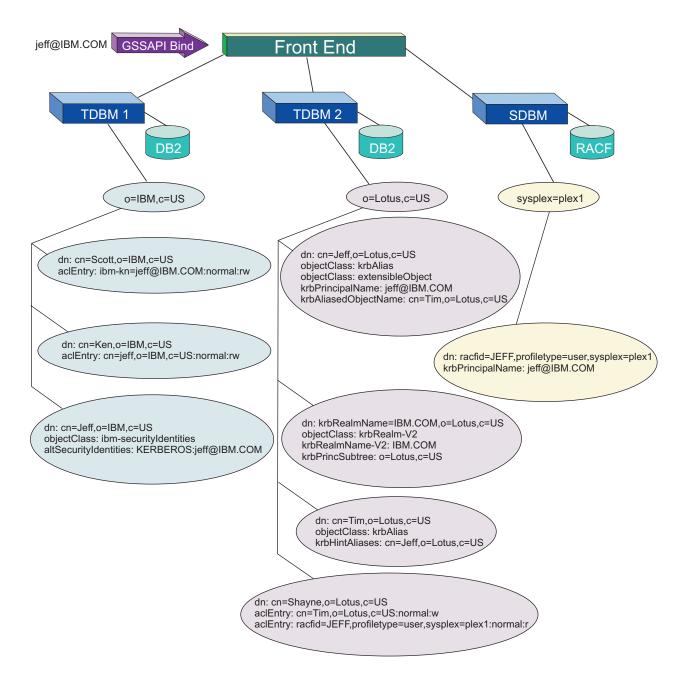


Figure 35. Kerberos directory example

Assume that Kerberos support has been enabled for this server, all backends have turned identity mapping on, and the JEFF user ID has performed a kinit to acquire a Kerberos ticket before issuing the GSS API Kerberos bind.

The user JEFF with a Kerberos identity of jeff@IBM.COM is performing a Kerberos GSS API Bind to an LDAP server which has been configured with two TDBM backends and an SDBM backend.

During the bind process the Kerberos identity jeff@IBM.COM by default is mapped to ibm-kn=jeff@IBM.COM and this value is added to the list of DNs that is used for access control.

After default mapping is performed, each of the backends attempt to perform identity mapping:

- 1. The TDBM 1 backend first looks for the Kerberos realm object with a krbRealmName-V2 = IBM.COM and will not find one. Now the backend attempts to find the entry that contains altSecurityIdentities = KERBEROS: jeff@ibm.com. The entry with the DN cn=Jeff,o=IBM,c=US matches this criteria and the DN is added to the alternate DN list.
- 2. Now the server moves on to the TDBM 2 backend and tries to find the Kerberos realm object with a krbRealmName-V2 = IBM.COM. This time the realm object is found so all of the krbPrincSubtree values of the realm object are collected. Next, the server searches each of these subtrees (in this example, only the o=Lotus.c=US subtree) for entries that contain krbPrincipalName = jeff@IBM.COM. In this backend the entry cn=Jeff,o=Lotus,c=US is found and is added to the DN list. Next the JEFF entry is checked for the krbAliasedObjectName attribute. There is a krbAliasedObjectName specified so authorization of the alias needs to be performed. The alias is cn=Tim,o=Lotus,c=US so the Tim entry must be checked for the attribute krbHintAliases with a value of cn=Jeff,o=Lotus,c=US. This value does exist so the DN cn=Tim,o=Lotus,c=US is added to the access control DN list.

Note: If the value cn=Jeff,o=Lotus,c=US did not exist in Tim's krbHintAliases, then Tim did not want you to alias him, so the DN cn=Tim,o=Lotus,c=US would not have been added to the DN list.

3. Finally, the server gets to the SDBM backend and invokes a RACF API that attempts to map the Kerberos identity jeff@IBM.COM to its associated RACF ID. In this example, the API returns the JEFF user ID and the DN racfid=JEFF, profiletype=user, sysplex=plex1 is constructed and added to the list of access control DNs.

At this point, the bind has completed and the list of DNs that is used for access control is as follows:

```
ibm-kn=jeff@IBM.COM
cn=Jeff,o=IBM,c=Us
cn=Jeff,o=Lotus,c=Us
cn=Tim.o=Lotus.c=US
racfid=JEFF,profiletype=user,sysplex=plex1
```

Group gathering can now be performed on the entire list of DNs.

Now that jeff@IBM.COM is bound to the server and his list of alternate DNs has been generated, he now has authority to perform other operations:

- Because jeff@IBM.COM was mapped to ibm-kn=jeff@IBM.COM he has read and write permission to normal data in the cn=Scott,o=IBM,c=US entry.
- jeff@IBM.COM also has read and write permission to the normal data in the cn=Ken,o=IBM,c=US entry due to his identity also being mapped to cn=Jeff,o=IBM,c=US.
- Modify operations would be permitted on the cn=Shayne,o=IBM,c=US entry since jeff@IBM.COM was also mapped to cn=Tim,o=Lotus,c=US and Tim has write access to Shayne.
- Read access is also permitted on the cn=Shayne,o=IBM,c=US entry because jeff@IBM.COM was mapped to the SDBM DN racfid=JEFF, profiletype=user, sysplex=plex1 who has read permission to the cn=Shayne,o=IBM,c=US entry.

You can see from this example that your access control is based on the combination of all the mapped DN's access control permissions.

Kerberos operating environments

Because Kerberos Version 5 is interoperable with other Kerberos 5 implementations, there are a variety of different operating environments that can exist.

- Another KDC other than the z/OS KDC can be used to store Kerberos principals. Users get a ticket from the other KDC rather than the z/OS KDC. The LDAP server could also be registered to this other KDC. However a trusted realm between the z/OS KDC and the external KDC must be established.
- Another KDC can be used along with the z/OS KDC where both will store Kerberos identities. In this scenario a trusted realm needs to be configured between the two realms.

• z/OS user IDs can be set up to contain an external KDC's Kerberos identities so when SDBM identity mapping is performed you can still be mapped to a RACF ID if you are strictly using the external or foreign KDC for Kerberos identities. This is done by setting up a KERBLINK and a trusted realm. The following KERBLINK example adds the foreign Kerberos identity jeff@KRB2000.IBM.COM to the RACF user JEFF:

RDEFINE KERBLINK /.../KRB2000.IBM.COM/jeff APPLDATA('JEFF')

For information about how to set up trusted realms and **KERBLINK**, refer to *z/OS: Security Server* Network Authentication Service Administration.

Chapter 20. Native authentication

LDAP has the ability to authenticate to the Security Server through TDBM by supplying a Security Server password on a simple bind to a TDBM backend. Authorization information is still gathered by the LDAP server based on the DN that performed the bind operation. The LDAP entry that contains the bind DN should contain either the **ibm-nativeld** attribute or **uid** attribute to specify the ID that is associated with this entry. Note that the SDBM backend does not have to be configured. The ID and password are passed to the Security Server and the verification of the password is performed by the Security Server. Another feature of native authentication is the ability to change your Security Server's password by issuing an LDAP modify command.

Note: The use of RACF passtickets is supported by the z/OS LDAP server. It is recommended that the LDAP server be run as a started task if RACF passticket support will be used. The job name associated with the LDAP server started task should be used as the application name when generating RACF passtickets. Refer to z/OS: Security Server RACF Macros and Interfaces for more information about RACF passtickets.

Initializing native authentication

To enable native authentication, perform the following steps:

- 1. Install and configure RACF or another Security Server.
- 2. Configure an LDAP server to run TDBM and start the server. Specify the native authentication options in your configuration file. For example:

```
TDBM Section
useNativeAuth SELECTED
"nativeAuthSubtree" o=IBM,c=US
"nativeAuthSubtree" o=Lotus,c=US
nativeUpdateAllowed YES
```

- 3. Load the native authentication related schema elements into the TDBM backend.
- 4. Be sure that the entries that are to perform native authentication contain either the **ibm-nativeld** attribute or a single-valued **uid** attribute with the appropriate Security Server ID as its value. It is important to note that a multi-valued **uid** without an **ibm-nativeld** causes the bind to fail because the LDAP server does not know which ID to use.

Updating the schema for native authentication

In order to enable native authentication, the directory schema must contain the native authentication related schema elements which are defined in **schema.IBM.Idif**. For more information on schema updates see "Updating the schema" on page 173.

Following is the native authentication attribute type:

ibm-nativeld

Allows you to specify the ID that is to be associated with this entry.

Following is the native authentication object class:

ibm-nativeAuthentication

Allows you to specify the **ibm-nativeld** attribute in entries.

Defining participation in native authentication

There are many different configuration options for native authentication which are discussed in this section.

The main configuration option, useNativeAuth, can be set to selected, all, or off. If you want all entries in a certain subtree to participate in native authentication then you would choose all for this option. However, if you would like specific entries in the specific subtrees to be subject to native authentication, then choose selected for the useNativeAuth option. When selected is used, only entries with the ibm-nativeId attribute will be subject to native authentication.

In order for an entry to bind natively or perform a native password modify, that entry must contain a mapping to the Security Server identity that is associated with the user. This can be accomplished by using either the ibm-nativeld attribute or the UID attribute that is defined in schema.user.ldif. If your directory entries already contain a single-valued **UID** attribute (which holds the Security Server user ID), then these entries are already configured for native authentication if you plan on using the useNativeAuth all option. If you do not plan on using UIDs for mapping, then you can specify the ibm-nativeld attribute for your Security Server ID associations and this attribute is used with selected or all specified for the useNativeAuth option. If both the ibm-nativeId and UID attributes exist in an entry, the ibm-nativeId value is used. The userid specified be either the UID or ibm-nativeld attributes must contain a valid OMVS segment in the Security Server.

If you use the useNativeAuth option, also specify the nativeUpdateAllowed option to enable native password changes in the Security Server to occur through the TDBM backend.

Next, consider what portions of your directory should have the ability to participate in native authentication. If the entire directory should participate, then set the nativeAuthSubtree configuration option to all. If there are different subtrees in your directory which contain entries that need to bind natively or perform native password modifications, then you need to list all the subtrees with the nativeAuthSubtree configuration option.

Note: If the DN that is listed in the nativeAuthSubtree option contains a space character in it, then the entire DN must be enclosed in quotes in the configuration file.

As mentioned above, there are two LDAP operation affected: bind and password modify. There is a set of criteria that is used to determine if an entry actually participates in native authentication. This criteria changes depending on the configuration options that have been selected. The following table outlines all the possible operating modes for native authentication.

Table 37. Operating modes for native authentication

Operation	useNativeAuth	nativeUpdate Allowed	ibm-nativeld	UID	Behavior
Bind	selected	any value	User1		Entry is configured correctly and native authentication is attempted.
Bind	selected	any value		User1	Entry is not correctly configured for native authentication so an LDAP simple bind is attempted. The UID attribute is not used when useNativeAuth is selected.
Bind	selected	any value			Entry has not been configured for native authentication so an LDAP simple bind is attempted.
Bind	all	any value	User1	User2	The ibm-nativeld attribute is used to attempt native authentication.
Bind	all	any value		User1	Entry is configured correctly and native authentication is attempted.

Table 37. Operating modes for native authentication (continued)

Operation	useNativeAuth	nativeUpdate Allowed	ibm-nativeld	UID	Behavior
Bind	all	any value			For ease of implementation, a LDAP simple bind is attempted, even though you have specified that all entries should use native authentication. Eventually this entry should be configured correctly.
Modify/ Replace (password)	selected	Yes	User1		Operation is not allowed because the entry is configured for native authentication. A modify/delete followed by a modify/add must be performed.
Modify/ Replace (password)	selected	Yes			Entry is not configured for native authentication so a regular LDAP password replace is attempted.
Modify/ Replace (password)	all	Yes			Operation is not allowed. modify/delete followed by a modify/add must be performed.
Modify/ Delete (password)	selected	Yes	User1		Entry is configured for native authentication so the value specified is used to change USER1 Security Server password if a modify/add follows this operation. If a modify/add does not follow, then the operation fails.
Modify/ Delete (password)	selected	Yes			Entry is not configured for native authentication so a regular LDAP modify/delete is attempted.
Modify/ Delete (password)	all	Yes	User1	User2	Entry is configured for native authentication so the value specified is used to change USER1 Security Server password if a modify/add follows this operation. If a modify/add does not follow, then the operation fails.
Modify/ Delete (password)	all	Yes		User1	Entry is configured for native authentication so the value specified is used to change USER1 Security Server password if a modify/add follows this operation. If a modify/add does not follow, then the operation fails.
Modify/ Delete (password)	all	Yes			A regular LDAP modify/delete will be tolerated in this case to allow for old LDAP passwords stored in TDBM to be removed.

Table 37. Operating modes for native authentication (continued)

Operation	useNativeAuth	nativeUpdate Allowed	ibm-nativeld	UID	Behavior
Modify/ Add (password)	selected	Yes	User1		If a password modify/delete was previously performed, then a Security Server password change for USER1 is attempted. If the modify/delete had not been performed, then the operation fails. Also, if the Security Server ID is not defined, the operation fails.
Modify/ Add (password)	selected	Yes			Entry is not configured for native authentication so a regular LDAP modify/add is attempted.
Modify/ Add	all	Yes	User1	User2	If a password modify/delete was previously performed then a Security Server password change for USER1 is attempted. If the modify/delete had not been performed, then the operation fails. Also, if the Security Server ID is not defined, the operation fails.
Modify/ Add	all	Yes		User1	If a password modify/delete was previously performed, then a Security Server password change for USER1 is attempted. If the modify/delete had not been performed, then the operation fails. Also, if the Security Server ID is not defined, the operation fails.
Modify/ Add	all	Yes			Operation fails because the entry is not correctly configured for native authentication.
Add (entry with password)	selected		User1		Entry is configured for native authentication so adding an entry with a password is not allowed.
Add (entry with password)	selected			User1	Entry is not configured for native authentication so the operation is attempted.
Add (entry with password)	selected				Entry is not configured for native authentication so the add operation is attempted.
Add (entry with password)	all		User1	User2	Operation fails. Native entries do not need LDAP passwords.
Add (entry with password)	all			User1	Operation fails. Native entries do not need LDAP passwords.
Add (entry with password)	all				Operation fails. Native entries do not need LDAP passwords.

Notes: This table assumes that the entry is located within native authentication subtrees. If a Security Server ID has been specified for a native entry but has not yet been defined in Security Server, the LDAP server will attempt an LDAP simple bind. If a native entry is configured with a multi-valued **UID** and no **ibm-nativeld**, then the operation fails because the LDAP server does not know which value to use as the user ID.

Updating native passwords

Performing a native password modify is as simple as issuing an Idapmodify command to perform a delete followed by an add of the userpassword attribute. Specify the current password on the delete statement followed by the new password on the add statement. The delete must occur before the add for native password modify. For example, if the file pw.mod contains:

cn=You,o=IBM,c=US -userpassword=oldpassword +userpassword=newpassword

then the following command modifies the native password (assuming the bind DN has the authority to do

ldapmodify ... -D cn=You,o=IBM,c=US -w oldpassword -f pw.mod

The following *errno* values returned by **password()** will have an LDAP reason code defined for them:

Table 38. The errno values returned by _passwd()

errno value	Reason	Text
EMVSERR	R004107	The password function failed; not loaded from a program controlled library.
EMVSSAF2ERR	R004108	TDBM backend password API resulted in an internal error.
EMVSEXPIRE	R004109	The password has expired.
EMVSSAFEXTRERR	R004110	The userid has been revoked.
EMVSPASSWORD	R004128	Native authentication password change failed. The new password is not valid or does not meet requirements.
EACCES	R004111	The password is not correct.
EINVAL	R004112	A bind argument is not valid.
ESRCH	R004118	Entry native user ID (ibm-nativeld,uid) is not defined to the Security Server.

If the Idapmodify command above fails with LDAP return code LDAP_INVALID_CREDENTIALS and LDAP reason code:

R004109 The password has expired.

then it is possible to change the RACF password of a TDBM entry participating in native authentication by doing an LDAP simple bind. The simple bind can occur as part of an LDAP function such as search, add, or modify. The password change is provided in the password portion of the LDAP simple bind. The password must be in the following format:

password/newpassword

The forward slash (/) is used as the indication of a password change during the LDAP simple bind. Password changes made using the LDAP simple bind to a TDBM entry participating in native authentication are subject to the system password rules. A password change will fail with LDAP return code LDAP_INVALID_CREDENTIALS and LDAP reason code of:

R004128 Native authentication password change failed: The new password is not valid, or does not meet requirements.

if the new password does not pass the rules established on the system. Note that once the bind succeeds, the password is changed even if the LDAP function eventually fails. Assuming TDBM entry cn=User1,ou=END,o=IBM,c=US is participating in native authentication, the following command changes the RACF password for user USER1 from abc to def: ldapsearch -h ldaphost -p ldapport -D "cn=User1,ou=END,o=IBM,c=US" -w abc/def -b "ou=END,o=IBM,c=US"\ "objectclass=*" Note: LDAP ACLs must be set properly to allow update of the userpassword attribute for the password modification to complete successfully. The distinguished name provided on the -D parameter of the **Idapmodify** command must have authority to update the *userpassword* attribute. To allow each individual user to update their own password, an LDAP ACL should be established to permit them to write userpassword attribute values. If you are running z/OS R2 or higher, or have applied OW50971 to OS/390 V2R10 or z/OS V1R1, you can use the special cn=this identity to establish the LDAP ACL. Run the following **Idapmodify** command to establish the LDAP ACL: ldapmodify -D admin DN -w admin DN password -f /tmp/aclmod.ldif where the file /tmp/aclmod.ldif looks like: dn:o=Your Company

You should substitute the root of your directory tree for the dn:o=Your Company line in the LDIF file. This will allow each user defined for Native Authentication to update their own RACF password via LDAP.

Example of setting up native authentication

aclEntry:access-id:cn=this:critical:rwsc

changetype:modify

aclPropagate:TRUE

add:x

The following diagram shows an example of how you could set up native authentication.

Note: Due to space limitations of the diagram, the entries in the example do not contain all of the necessary information to make them valid directory entries. For example, object classes and required attributes have been left out of many of the entries.

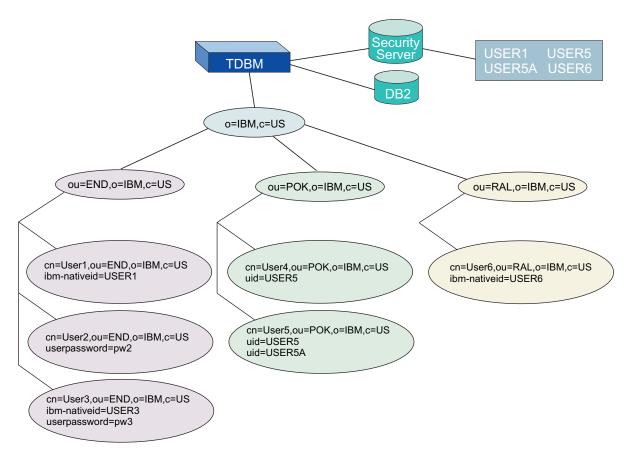


Figure 36. Native authentication example

Example 1

· Assuming these settings:

useNativeAuth selected nativeUpdateAllowed yes "nativeAuthSubtree" ou=END,o=IBM,c=US "nativeAuthSubtree" ou=POK,o=IBM,c=US

the following table indicates the results of operations involving each user entry:

Table 39. Behavior of native authentication in example 1

LDAP entry	Operation	Behavior
cn=User1,ou=END,o=IBM,c=US	Bind	Can bind natively because the entry contains a valid ibm-nativeld .
	Native password change	Can change this native password because the entry contains a valid ibm-nativeld .
	Modify/replace (user password)	Cannot perform a modify/replace of user password because the entry is subject to native authentication and password replace is not allowed.
cn=USER2,ou=END,o=IBM,c=US	All	Entry is not configured for native authentication so all operations are regular LDAP operations.
cn=USER3,ou=END,o=IBM,c=US	Bind	Attempts native authentication but because the Security Server ID is not defined, a regular LDAP bind is performed.

Table 39. Behavior of native authentication in example 1 (continued)

LDAP entry	Operation	Behavior
	Native password change	Native password change is attempted but eventually fails because the Security Server ID USER3 is not defined.
	Modify/replace (user password)	An attempt to modify/replace the user password attribute will fail because the entry is configured for native authentication.
cn=USER4,ou=POK,o=IBM,c=US	All	Performs regular LDAP operations because the entry does not contain the ibm-nativeld attribute.
cn=USER5,ou=POK,o=IBM,c=US	All	Performs regular LDAP operations because the entry does not contain the ibm-nativeld attribute.
cn=USER6,ou=RAL,o=IBM,c=US	All	Performs regular LDAP operations because the entry does not exist in a native subtree.

Example 2

· Now, assuming these settings:

useNativeAuth all nativeUpdatAllowed yes "nativeAuthSubtree" ou=END,o=IBM,c=US "nativeAuthSubtree" ou=POK,o=IBM,c=US

the following table indicates the results of operations involving each user entry:

Table 40. Behavior of native authentication in example 2

LDAP Entry	Operation	Behavior
cn=User1,ou=END,o=IBM,c=US	Bind	Can bind natively because the entry contains a valid ibm-nativeld .
	Native password change	Can change this native password because the entry contains a valid ibm-nativeld .
	Modify/replace (user password)	Cannot perform a modify/replace of user password because the entry is subject to native authentication and password replace is not allowed.
cn=USER2,ou=END,o=IBM,c=US	Bind	Because there are no native attributes in this entry, a regular LDAP bind is attempted.
	Native password change	An attempt to change the user password attribute fails because the entry is not correctly configured for native authentication. The entry should contain either the ibm-nativeld attribute or UID attribute.
	Modify/replace (user password)	Modify/replace of user password is not allowed because useNativeAuth is all .
cn=USER3,ou=END,o=IBM,c=US	Bind	Attempts native authentication but because the Security Server ID is not defined, a regular LDAP bind is performed.
	Native password change	Native password changes are attempted but eventually fail because the Security Server ID is not defined.
	Modify/replace (user password)	An attempt to modify/replace the user password attribute will fail because the entry is configured for native authentication.

Table 40. Behavior of native authentication in example 2 (continued)

LDAP Entry	Operation	Behavior
cn=USER4,ou=POK,o=IBM,c=US	Bind	Can bind natively because the entry contains the UID attribute.
	Native password change	Can change the native password because the entry contains a valid UID .
	Modify/replace (user password)	An attempt to modify/replace the user password attribute will fail because the entry is configured for native authentication.
cn=USER5,ou=POK,o=IBM,c=US	Bind	Native bind fails because 2 UID values exist.
	Native password change	Change password fails because 2 UID values exist.
	Modify/replace (user password)	Modify/replace of user password attribute is not allowed because useNativeAuth is all .
cn=USER6,ou=RAL,o=IBM,c=US	All	Performs regular LDAP operations because the entry does not exist in a native subtree.

Using native authentication with Web servers

Many Web servers provide a user ID and password challenge for authentication. These can take advantage of native authentication. The Web server must be configured to do LDAP authentication. When the challenge to do LDAP authentication is presented, the user can enter the Security Server user ID and password (from the system where the LDAP server is running). The Web server will search the LDAP directory for an entry where UID equals the input user ID. The Web server will use the returned DN and the inputted password to do an Idap_simple_bind(). When the LDAP server determines this entry is subject to native authentication, it will retrieve the ibm-nativeld or UID value and verify the password with the Security Server. Note that if useNativeAuth is set to selected, it may be necessary to place the Security Server user ID into both the UID and ibm-nativeld attributes of this entry to allow the Web server processing to work correctly with native authentication.

Chapter 21. CRAM-MD5 and DIGEST-MD5 Authentication

The z/OS LDAP server allows clients to authenticate to the server by using CRAM-MD5 (Challenge Response Authentication Mechanism - RFC 2104) and DIGEST-MD5 (RFC 2831) SASL bind mechanisms. CRAM-MD5 and DIGEST-MD5 bind mechanisms are multi-stage binds where the server sends the client a challenge and then the client sends a response back to the server to complete authentication. The client response contains a hash, which is encoded according to the specifications of the DIGEST-MD5 or CRAM-MD5 RFC, and a username. The username that is specified on the client is also known as the authentication identity, which is used to perform the authentication with the server. On the z/OS LDAP server, the username is the **uid** attribute value of an entry that is targeted for authentication. When the server gets the response from the client, the response is parsed and the server calculates its own hash with the password found for the **uid** attribute value in the backend. If the server hash is equal to the client hash, then the client and the server are authenticated.

DIGEST-MD5 and CRAM-MD5 SASL bind mechanisms are much better than simple binds since they do not pass the credentials in the clear, which are more liable to interception by snoopers. Also, CRAM-MD5 and DIGEST-MD5 bind mechanisms on the z/OS LDAP server do not require any additional products to be installed or configured in order to have this added functionality.

DIGEST-MD5 authentication is much stronger than CRAM-MD5 authentication because it prevents chosen plaintext password attacks. During a DIGEST-MD5 authentication, there is additional information that is passed between the server and the client during the challenges and responses than compared to the CRAM-MD5 algorithm. Thus, it is more difficult to use the DIGEST-MD5 hashing algorithm to decipher the password for the authentication identity than the CRAM-MD5 hashing algorithm.

DIGEST-MD5 restrictions on the z/OS LDAP server:

- 1. The options for integrity and encryption protection on DIGEST-MD5 binds are not supported.
- 2. The unspecified userid form of the authorization identity is not supported; however, the DN version is supported on the z/OS LDAP client and server.
- 3. An authorization DN is used to obtain the authority of another LDAP user after successfully authenticating based upon the authentication identity. The z/OS LDAP server does not support specifying an authorization DN that is different than the authentication identity's DN.
- 4. Subsequent authentication is not supported.

Considerations for setting up a TDBM backend for CRAM-MD5 and DIGEST-MD5 Authentication

The following are considerations for setting up a TDBM backend for CRAM-MD5 and DIGEST-MD5 authentication:

1. In order to use CRAM-MD5 and DIGEST-MD5 authentication on the z/OS LDAP server, the entries that you bind with must contain a uid attribute value. The uid attribute value is used by the z/OS LDAP server as the authentication identity or username when attempting a CRAM-MD5 or DIGEST-MD5 authentication bind. The uid attribute is found in the schema file: /usr/lpp/ldap/etc/schema.user.ldif. It is strongly suggested that the uid attribute values be unique across every TDBM backend that is configured on the LDAP server. When attempting a CRAM-MD5 or DIGEST-MD5 bind, the username that is specified on the z/OS LDAP operation utilities must correspond to the uid attribute value of the entry that you want to authenticate as. A search of every TDBM backend that is configured on the LDAP server will be done to look for the uid attribute or username value that was entered on the client. If this search finds more than one DN entry that has the same uid value, then the bind will not be successful because it is not known which entry is targeted for authentication. However, this problem can be avoided by specifying an authorization DN. If the authentication identity or uid attribute value is present in the authorization DN's entry, then the bind will be successful assuming that the password is correct; otherwise it will fail.

243

- In order for CRAM-MD5 and DIGEST-MD5 binds to work properly, the userpassword attribute for the
 entry must be in either clear text (not recommended) or in the DES 2-way encryption algorithm
 (recommended). See Chapter 7, "Preparing the backends, SSL/TLS, and password encryption" on
 page 39.
- 3. CRAM-MD5 and DIGEST-MD5 binds are not supported with entries that are participating in Native Authentication.
- 4. CRAM-MD5 and DIGEST-MD5 binds are not supported to the SDBM backend.

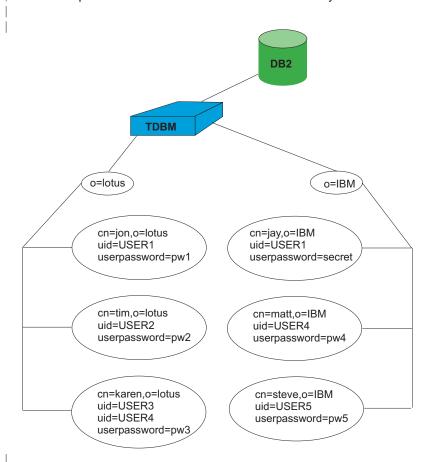
CRAM-MD5 and DIGEST-MD5 configuration parameter

The **digestRealm** optional configuration parameter allows for the specification of a realm name to be used to help create the CRAM-MD5 and DIGEST-MD5 hashes. The value of this parameter gets passed on the initial challenge from the server to the client once it has been decided that a CRAM-MD5 or DIGEST-MD5 bind is desired. See the **digestRealm** parameter in the "Configuration file options" on page 53. This parameter defaults to the fully qualified DNS hostname of the LDAP server.

Example of setting up for CRAM-MD5 and DIGEST-MD5

The following diagram shows an example of how you could set up your entries in your TDBM backend(s).

Note: Due to space limitations of the diagram, the entries in the example do not contain all of the necessary information to make them valid directory entries. For example, object classes and required attributes have been left out of many of the entries.



The following table outlines what would happen if you attempt to do a CRAM-MD5 or DIGEST-MD5 bind from a client. The authentication identity or username refers to the **-U** option on the z/OS LDAP operation utilities, while the authorization DN is the **-D** option on the z/OS LDAP operation utilities. See z/OS:

Security Server LDAP Client Programming for more details on the LDAP operation utilities. This table assumes that native authentication is not turned on under the subtrees: o=lotus and o=IBM.

Table 41.

Authentication Identity	Authorization DN	Password	Behavior
USER2		pw2	Bind is successful to cn=tim,o=lotus
USER2	cn=tim,o=lotus	pw2	Bind is successful to cn=tim,o=lotus
USER2	cn=jon,o=lotus	pw2	Bind is not successful because the authorization DN (cn=jon,o=lotus) does not equal the authentication DN (cn=tim,o=lotus)
USER1		pw1	Bind is not successful, because there are multiple entries with the same username/uid value: cn=jon,o=lotus and cn=jay,o=IBM
USER1	cn=jay,o=IBM	secret	Bind is successful to cn=jay,o=IBM because the authentication DN (cn=jay,o=IBM) equals the authorization DN (cn=jay,o=IBM)
USER1	cn=jon,o=lotus	pw1	Bind is successful to cn=jon,o=lotus because the authentication DN (cn=jon,o=lotus) equals the authorization DN (cn=jon,o=lotus)
USER3		pw3	Bind is successful to cn=karen,o=lotus
USER4	cn=karen,o=lotus	pw3	Bind is successful to cn=karen,o=lotus
USER3	cn=karen,o=lotus	bad	Bind is not successful to authentication DN (cn=karen,o=lotus) and authorization DN (cn=karen,o=lotus) because the password is incorrect.
USER6	cn=notheree,o=lotus	pw6	Bind is not successful because the authentication DN (cn=matt,o=IBM) does not equal the non-existent authorization DN (cn=nothere,o=lotus)
BAD		pw1	Bind is not successful because a uid value equal to BAD was not found in any of the entries in the TDBM backend.

Chapter 22. Using extended operations to access Policy Director data

The extended operations (EXOP) backend supports two extended operations that open a connection to the target LDAP server to access z/OS Policy Director data. The **IbmLDAPProxyControl** determines the target LDAP server. To set the target LDAP server when using z/OS Policy Director, use the RACF PROXY segment. See *z/OS: RACF Security Administrator's Guide* for more information.

The LDAP extended operations are **GetDnForUserid** and **GetPrivileges**. These extended operations are generated when an application on z/OS calls the AZN APIs (refer to *Policy Director Authorization Service for z/OS and OS/390: Customization and Use* for more information on using the AZN APIs). When the EXOP backend receives a request for either of these two operations, it uses the required **IbmLDAPProxyControl** to open an LDAP connection to a target LDAP server that has been set up to store Policy Director data. Then, depending on the request, the EXOP backend issues LDAP requests to the target server to retrieve the appropriate data.

GetDnForUserid extended operation

For the **GetDnForUserid** extended operation, the EXOP backend retrieves all of a user ID's distinguished names (DNs) stored in the target LDAP server. The client can filter the DNs returned by the EXOP backend by specifying a search base and object class names. The sequence of events for this extended operation is:

- If the client does not specify a search base, the EXOP backend searches for the DN of all entries in all of the target server's naming contexts that contain an **ibm-nativeid** attribute set to the specified user ID and whose set of object classes include all of the optional specified object classes.
- If the EXOP backend does not receive entries from the target LDAP server for this first set of searches, it attempts a similar set of searches, maintaining the filtering based on the optional object classes. For the second set of searches, however, instead of searching for entries with an ibm-nativeid attribute set to the specified user ID, it searches for entries with a uid attribute set to the specified user ID.

If the client does specify a search base, the EXOP backend will attempt the same sequence of searches described above, but instead of searching all of the target LDAP server's naming contexts, it only searches the naming context specified in the search base.

Appendix H, "Supported extended operations" on page 435 summarizes some different error scenarios for this extended operation and the EXOP backend's response to such scenarios.

GetPrivileges extended operation

For the **GetPrivileges** extended operation, the EXOP backend retrieves all of a subject's Policy Director data. This subject is specified by its DN. The client can specify an optional domain name if the subject does not exist in the domain named DEFAULT. Refer to "GetPrivileges" on page 436 for an ASN.1 description of all of the data that the EXOP backend retrieves when it receives this extended operations request.

To satisfy this request, the EXOP backend performs many searches then combines all of the results prior to returning it to the client. Furthermore, some of the searches may require searches across all of the target LDAP server's naming contexts. For example, to find the groups the subject is a member of, the EXOP backend performs searches under all of the target LDAP server's naming contexts.

Appendix H, "Supported extended operations" on page 435 summarizes some different error scenarios for this extended operation and the EXOP backend's response to such scenarios.

These extended operations are used by z/OS Policy Director. More information about z/OS Policy Director is in Policy Director Authorization Service for z/OS and OS/390: Customization and Use.

Chapter 23. Using access control

Access control of information in the LDAP server is specified by setting up Access Control Lists (ACLs). ACLs provide a means to protect information stored in an LDAP directory. Administrators use ACLs to restrict access to different portions of the directory, or specific directory entries. LDAP directory entries are related to each other by a hierarchical tree structure. Each directory entry (or object), contains the entry's distinguished name, a set of attributes and their corresponding values. When using the TDBM backend, ACLs are created and managed using the **Idap_modify** APIs. ACLs can also be entered using the **Idif2tdbm** utility.

ACLs are represented by a set of attributes which appear to be a part of the entry. The attributes associated with access control, such as **entryOwner**, **ownerPropagate**, **aclEntry** and **aclPropagate** are unusual in that they are logically associated with each entry, but can have values which depend upon other entries higher in the directory hierarchy. Depending upon how they are established, these attribute values can be explicit to an entry, or inherited from an ancestor entry.

Use of LDAP's SDBM backend allows a user to be authenticated to the directory namespace using the RACF ID and password. The RACF identity becomes associated with the user's RACF-style distinguished name that was used on the LDAP bind operation. It is then possible to set up ACLs for entries managed by the TDBM backend using RACF-style user and group DNs. This controls access to TDBM database directory entries using the RACF user or group identities.

Access control attributes

Access to LDAP directory entries and attributes is defined by Access Control Lists (ACLs). Each entry in the directory contains a special set of attribute/value pairs which describe who is allowed to access information within that entry. Table 42 shows the set of attributes which are related to access control. More in-depth information about each attribute is given following the table.

Starting in z/OS V1R4, it is possible to specify access control settings for individual attribute types. This is called attribute-level access control. Also, added in z/OS V1R4 is the ability to explicitly deny access to information. In prior releases, all ACL information indicated what access was granted with the implication that access to all other information was denied.

Table 42. TDBM ACL and entry owner attributes

ACL attributes	
aclEntry	This is a multi-valued attribute that contains the names and permissions associated with those names that have access to information in the directory entry (or the entry along with the subtree of information below the entry, depending on the setting of the aclPropagate attribute).
aciPropagate	This is a single-valued boolean attribute which indicates whether the aclEntry information applies only to the directory entry it is associated with or to the entire subtree of information including and below the directory entry it is associated with. Note that propagation does not apply to entries that have an explicit aclEntry defined for the entry and that propagation stops at the next propagating ACL (aclPropagate=TRUE) that is encountered in the directory subtree.
aciSource	This is a single-valued attribute that is not modifiable. This attribute is managed by the directory server and cannot be changed by the Idapmodify command. This attribute, accessible for any directory entry, indicates the distinguished name of the entry that holds the ACL that applies to the entry. This attribute is useful in determining which propagating ACL is used to control access to information in the directory entry.
Entry owner attributes	

Table 42. TDBM ACL and entry owner attributes (continued)

entryOwner	This is a multi-valued attribute that contains the distinguished names of users or groups that are considered owners of the directory entry (or the entry along with the subtree of information below the entry, depending on the setting of the ownerPropagate attribute).
ownerPropagate	This is a single-valued boolean attribute which indicates whether the entryOwner information applies only to the directory entry it is associated with or to the entire subtree of information including and below the directory entry it is associated with. Note that propagation does not apply to entries that have an explicit entryOwner (ownerPropagate=FALSE) defined for the entry and that propagation stops at the next propagating entryOwner that is encountered in the directory subtree.
ownerSource	This is a single-valued attribute that is not modifiable. This attribute is managed by the directory server and cannot be changed by the Idapmodify command. This attribute indicates the distinguished name of the entry that holds the entryOwner that applies to the entry. This attribute is useful in determining which propagating entryOwner is used to control access to information in the directory entry.

aclEntry attribute

An aclEntry is a multi-valued attribute which contains information pertaining to the access allowed to the entry and each of its attributes. An aclEntry lists the following types of information:

- · Who has rights to the entry (scope of the protection). Also called the subject.
- What attributes, and classes of attributes (attribute access classes) that the subject has access to.
- · What rights the subject has (permissions to attribute and classes of attributes).

Scope of protection

The scope of the protection is based on the following three types of privilege attributes:

access-id

The distinguished name of an entry being granted access.

group The distinguished name of the group entry being granted access.

role The distinguished name of the group entry being granted access. (Use this for the TDBM backend only.)

Access control groups have an object class of accessGroup. The accessGroup object class is a subclass of the groupOfNames object class. Groups identified in an aclEntry attribute value must be of object class accessGroup.

Privilege attributes take the form of type:name where type refers to either access-id, group, or role and name is the distinguished name.

Note: The distinguished name that is used need not be the name of an entry in the directory. The distinguished name is the name that represents the user that has authenticated to the directory server.

The *type*: portion of this clause is optional.

Examples

In this example, the DN type is access-id and the DN itself is cn=personA, ou=deptXYZ, o=IBM, c=US. access-id:cn=personA, ou=deptXYZ, o=IBM, c=US

In this example, the DN type is group and the DN itself is cn=deptXYZRegs, o=IBM, c=US. group:cn=deptXYZRegs, o=IBM, c=US

This is an example of how to use the RACF identity established with SDBM in a TDBM ACL.

```
access-id:racfid=YourID.profileType=user.sysplex=YourSysplex
group:racfid=YourGroup,profileType=group,sysplex=YourSysplex
```

Attribute access classes

Attributes requiring similar permission for access are grouped together in classes. Attributes are assigned to an attribute access class within the schema definitions. The **ibmAttributeTypes** attribute of the TDBM schema entry holds the attribute type's access class. The three attribute access classes are:

- normal
- sensitive
- critical

Each of these attribute access classes is discrete. If a user has write permission to **sensitive** attributes, then the user does not automatically have write permission to normal attributes. This permission must be explicitly defined.

The default attribute access class for an attribute is **normal**. By default, all users have read access to normal attributes. There are two additional attribute access classes used internally by LDAP for system attributes, but users are not permitted to define attributes in these classes. These attribute access classes are **restricted** and **system**. You can specify these access classes when granting permissions in ACLs.

For example, a person's name would typically be defined in the normal class. Perhaps a social security number would be considered sensitive, and any password information for the user would be considered critical. Following are some example definitions excerpted from schema.user.ldif. Note that the attribute userPassword is defined with access class critical.

```
attributetypes: (
 2.5.4.49
 NAME ( 'dn' 'distinguishedName' )
  DESC 'This attribute type is not used as the name of the object itself,
  EQUALITY distinguishedNameMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12
 USAGE userApplications
ibmattributetypes: (
 2.5.4.49
 ACCESS-CLASS normal
attributetypes: (
  2.5.4.35
 NAME 'userPassword'
  DESC 'Holds a password value for a distinguished name.'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.5
  USAGE userApplications
ibmattributetypes: (
 2.5.4.35
 ACCESS-CLASS critical
```

Beginning in z/OS V1R4, it is possible to specify access controls on individual attributes. However, when defining schema an access class is always defined for the attribute type. If not specified, that attribute type is defined to belong to the normal class.

Access permissions

Following is the set of access permissions.

Table 43. Permissions which apply to an entire entry

Add	Add an entry below this entry
Delete	Delete this entry

Table 44. Permissions which apply to attribue access classes

Read	Read attribute values
Write	Write attribute values
Search	Search filter can contain attribute type
Compare	Compare attribute values

Syntax

Following is the **aclEntry** attribute value syntax:

```
[access-id:|group:|role:]<subject_DN>[granted_rights]
```

The subject_DN is any valid DN which represents the object (entry) to which privileges are being granted. The DN ends when the first granted rights keyword is detected.

The granted_rights is specified as follows where object_rights_list is one or more elements of the set [a|d], and attr_rights_list is one or more elements of the set [r|w|s|c].

```
[:object:[grant:|deny:]object rights list] [:normal:[grant:|deny:]attr rights list]
 [:sensitive:[grant:|deny:]attr_rights_list] [:critical:[grant:|deny:]attr_rights_list]
[:restricted:[grant:|deny:]attr_rights_list] [:system:[grant:|deny:]attr_rights_list]
 [:at.attr_name:[grant:|deny:]attr_rights_list]
```

Note: Restricted attributes are aclEntry, aclPropagate, entryOwner, and ownerPropagate. In order to update access control information. You must have permissions to read and write these attributes. The system attributes include aclSource and ownerSource and other attributes for which the server controls the values. In order to update access control information, you must have permission to read and write these attributes.

Following are some examples of valid aclEntry values:

```
access-id:cn=Tim, o=Your Company:normal:rwsc:sensitive:rsc:object:ad
role:cn=roleGroup, o=Your Company:object:ad:normal:rsc:sensitive:rsc
group:cn=group1, o=Your Company:system:csr:normal:sw
cn=Lisa, o=Your Company:normal:rwsc:sensitive:rwsc:critical:rwsc:restricted:rwsc:system:rwsc
cn=Ken, o=Your Company:normal:rsc
group:cn=group2,dc=yourcompany,dc=com:normal:rwsc:at.cn:deny:w:sensitive:grant:rsc
cn=Karen,dc=yourcompany,dc=com:at.cn:grant:rwsc:normal:deny:rwsc
cn=Mary,dc=yourcompany,dc=com:normal:rwsc:sensitive:rwsc:critical:deny:rwsc:at.userpassword:w
```

The TDBM access control implementation supports several "pseudo-DNs". These are used to refer to large numbers of subject DNs which, at bind time, share a common characteristic in relation to either the operation being performed or the target object on which the operation is being performed. Currently, three pseudo DNs are defined:

```
group: cn=anybody
group:cn=authenticated
access-id:cn=this
```

The group: cn=anybody refers to all subjects, including those that are unauthenticated (considered anonymous users). All users belong to this group automatically. The group: cn=authenticated refers to any DN which has been authenticated to the directory. The method of authentication is not considered. The access-id:cn=this refers to the bind DN which matches the target object's DN on which the operation is performed.

With the support in z/OS V1R4 for attribute-level permissions as well as grant/deny support, the order of evaluation of the separate permissions clauses is important. The access control permissions clauses are evaluated in a precedence order, not in the order in which they are found in the ACL entry value. With the new support, there are four types of permissions settings: access-class grant permissions, access-class deny permissions, attribute-level grant permissions, and attribute-level deny permissions. The precedence for these types of permissions is as follows (from highest precedence to lowest):

- attribute-level deny permissions
- · attribute-level grant permissions
- access-class deny permissions
- access-class grant permissions

Using this precedence, a deny permission takes precedence over a grant permission (for the same item specified) while attribute-level permissions take precedence over access-class permissions.

When multiple ACL entry values apply for a user, all of the permissions from all applicable ACL entry values are combined and then the precedence rules are applied. With this approach, it is possible for a user that is a member of two groups to result in having access to less information than if they were a member of only one of the groups.

See the "Access determination" on page 254 section for some examples will be used to illustrate this processing.

aclPropagate attribute

Each entry with an explicit ACL has associated with it an aclPropagate attribute. By default, the entry's explicit ACL is inherited down the hierarchy tree, and its owner propagate attribute is set to TRUE. If set to FALSE, the explicit ACL for the entry becomes an override, pertaining only to the particular entry. The aclPropagate syntax is Boolean. See "Propagating ACLs" on page 256 for more information.

aclSource attribute

Each entry has an associated aclSource. This reflects the DN with which the ACL is associated. This attribute is kept and managed by the server, but may be retrieved for administrative purposes. This attribute cannot be set, only retrieved.

In order to aid in migration from RDBM databases to TDBM databases, LDIF files can contain aclSource values for certain entries. The rule is as follows:

If aclSource is specified, the value must be the same as the distinguished name of the entry in which it is specified. If aclSource is specified and has a different value than the distinguished name of the entry within which it appears, the add and load of that entry will fail.

The derivation of aclSource is further explained in "Propagating ACLs" on page 256.

entryOwner attribute

Each entry has an associated entryOwner. The entryOwner might be a user or a group, similar to what is allowed within the aclEntry. However, the entryOwner subject has certain privileges over the entry.

Entry owners are, in essence, the administrators for a particular entry. They have full access on that particular entry, similar to the administrator DN. Note that the administrator DN has full permission on every entry in the database. The entryOwner attribute syntax is distinguishedName. However, for compatibility with previous releases, the distinguished name can be preceded with access-id:, group:, or role:.

Note: The distinguished name that is used need not be the name of an entry in the directory. The distinguished name is the name that represents the user that has authenticated to the directory server.

Entry owners are not constrained by permissions given in the aclEntry. They have complete access to any entry attribute, and can add and delete entries as desired.

Entry owners, the administrator DN, and users who have write permission for restricted attributes are the only people who are allowed to change the attributes related to access control.

ownerPropagate attribute

Owner propagation works exactly the same as ACL propagation. By default, owners are inherited down the hierarchy tree, and their owner propagate attribute is set to TRUE. If set to FALSE, the owner becomes an override, pertaining only to the particular entry. The ownerPropagate syntax is boolean.

ownerSource attribute

Each entry also has an associated ownerSource. This reflects the DN with which the owner values are associated. This attribute is kept and managed by the server, but can be retrieved for administrative purposes. This attribute cannot be set, only retrieved.

In order to support existing ACL support from the RDBM backend and aid in migration from RDBM databases to TDBM databases. LDIF files can contain **ownerSource** values for certain entries. The rule is as follows:

If **ownerSource** is specified, the value must be the same as the distinguished name of the entry in which it is specified. If **ownerSource** is specified and has a different value than the distinguished name of the entry within which it appears, the add and load of that entry will fail.

Initializing ACLs with TDBM

In z/OS LDAP V1R2 and later, the TDBM backend creates a default ACL entry in the "suffix entry" if one is not specified during the add of this entry (whether the add was done using Idapadd or Idif2tdbm). This improves performance of future ACL modifications made to an ACL placed on the suffix entry. The default ACL that is used is:

aclEntry: cn=anybody:normal:rsc:system:rsc aclPropagate: TRUE

Access determination

Each of the LDAP access permissions is discrete. One permission does not imply another permission. If a user (access-id) is given an aclEntry for a particular entry, this is the permission set the user receives. If no permission is given to the user, the user receives permission based on the entry's effective ACL (which might contain permissions for "anonymous" users), and the user receives the combined permissions of all listed access groups of which they are a member.

If the user is not listed on the ACL, and is not a member of any listed access group, then the user receives the permissions listed under the group:cn=Authenticated or group:cn=Anybody ACL entry, depending on whether the user was authenticated to the directory with an LDAP bind operation. If one or both ACL entries do not exist, access to the entry is completely denied.

Following are examples for permissions:

Example 1

aclentry: group:cn=Anybody:normal:rsc

In this example, unauthenticated (anonymous) users have permission to read, search and compare all attributes within the normal attribute access class. ACL entry values for unauthenticated users use pseudoDN cn=Anybody.

Example 2

```
aclentry: access-id:cn=personA,ou=deptXYZ,o=IBM,c=US:object:ad:normal:rwsc:sensitive:rwsc:critical:rsc
```

In this example, the user corresponding to cn=personA, ou=deptXYZ, o=IBM, c=US has permission to add entries below the entry, to delete the entry, to read, write, search and compare both normal and sensitive attributes, and to read, search and compare critical attributes.

Example 3

```
aclentry: group:cn=Authenticated:normal:rwsc:sensitive:rwsc
```

In this example, users who have authenticated to the directory where a specific aclEntry value does not apply, will be allowed to read, write, search, and compare, normal and sensitive attributes in the directory entry.

Example 4

```
aclEntry: cn=Tim,dc=yourcompany,dc=com:at.cn:deny:w:normal:rwsc
```

In this example, cn=Tim,dc=yourcompany,dc=com is granted read, write, search, and compare to **normal** attributes except for the **cn** attribute in which write access is denied. Note that the following ACL entry results in the same access:

```
aclEntry: cn=Tim,dc=yourcompany,dc=com:normal:rwsc:at.cn:deny:w
```

The evaluation of the permissions clauses is based on precedence, not order in the ACL entry value(s).

Example 5

```
aclEntry: cn=Karen,dc=yourcompany,dc=com:normal:rwsc:sensitive:rsc:at.userpassword:w:
critical:deny:rwsc
```

In this example, cn=Karen, dc=yourcompany, dc=com is granted read, search, and compare to normal and sensitive attributes, and write to normal attributes and the userpassword attribute. All access to critical attributes (except for write in userpassword) is turned off.

Example 6

```
aclEntry: group:cn=group1,dc=yourcompany,dc=com:normal:rwsc
aclEntry: group:cn=group2,dc=yourcompany,dc=com:sensitive:rwsc:at.cn:deny:w
```

In this example, a member of group1 only would be granted read, write, search, and compare to normal attributes. A member of both group1 and group2 would be granted read, write, search, and compare to **normal** and **sensitive** attributes, excluding write access to the **cn** attribute. This is an example where a member of both groups is granted access to less information than what is granted to each of the two groups individually.

| Example 7

```
aclEntry: access-id:cn=Tim,dc=yourcompany,dc=com:normal:rwsc:at.cn:rsc
```

In this example, cn=Tim,dc=yourcompany,dc=com is granted read, write, search, and compare on **normal** attributes and read, search, and compare on the **commonName** attribute. Note that

cn=Tim,dc=yourcompany,dc=com will also have write access to the **commonName** attribute by virtue of commonName having an access class of normal.

Example 8

aclEntry: access-id:cn=Tim,dc=yourcompany,dc=com:normal:rwsc:at.cn:deny:rsc

In this example, cn=Tim,dc=yourcompany,dc=com is granted read, write, search, and compare on **normal** attributes and denied read, search, and compare on the commonName attribute. Note that cn=Tim, dc=yourcompany, dc=com will still have write access to the commonName attribute by virtue of commonName having an access class of normal.

Attribute classes and searching

In order to read information from the directory, the user must have read permission for the attribute.

Filter

In order to use an attribute in a search filter supplied on a search operation, the user must have search permission for the attribute.

Compare

In order to perform a compare operation on an attribute/value combination, the user must have compare permission on the attribute.

Requested attributes

If the user has the search permission on all attributes contained in the filter, the server returns as much information as possible. All requested attributes that the user has read permission for are returned.

For example, let the aclEntry be

group:cn=Anybody:normal:rsc:sensitive:c:critical:c

and let a client perform an anonymous search

ldapsearch -b "c=US" "cn=LastName" title userpassword telephoneNumber

where title is a normal attribute, telephoneNumber is a sensitive attribute, and userpassword is a critical attribute. See the z/OS: Security Server LDAP Client Programming for more information about the Idapsearch utility.

Users performing anonymous searches are given the permission granted to the cn=Anybody group. In this example, permission exists to the filter since cn is in the normal attribute access class, and **cn=Anybody** has **s** (search) permission to the **normal** attribute access class. What is returned however, is only the title for any matching entry. The telephoneNumber and userPassword are not returned since cn=Anybody does not have read permissions on the sensitive and critical attribute access classes.

Propagating ACLs

ACLs can be set on any entry in the hierarchy. LDAP ACLs can propagate down through the directory hierarchy. These ACLs, called propagating ACLs, have the aclPropagate attribute set to TRUE. All descendents of this entry will inherit the ACL set at that point, unless overridden. In order to specify an ACL different from that of its parent, a new ACL must be explicitly set.

When setting the new ACL, there is again a choice of whether to propagate the ACL. If set to **TRUE**, the ACL will propagate down to all descendants. If set to FALSE, the ACL is not propagated; it instead becomes an override ACL. The ACL is not propagated down through the hierarchy, but instead applies only to the one particular entry that it is associated with within the hierarchy. If unspecified, aclPropagate is set to TRUE.

An entry without an explicit ACL receives its ACL from the nearest propagating ancestor ACL. Propagated ACLs do not accumulate as the depth in the tree increases. The scope of a propagated ACL is from the explicitly-set propagating ACL down through the tree until another explicitly-set propagating ACL is found.

Example of propagation

Following is the explicit ACL for entry ou=deptXYZ, o=IBM, c=US:

```
aclPropagate: TRUE
aclEntry: group:cn=deptXYZRegs, o=IBM, c=US:normal:rcs:sensitive:rsc
aclEntry: access-id:cn=personA, ou=deptXYZ, o=IBM, c=US:object:ad:normal:rwsc:sensitive:rwsc:critical:rsc
aclEntry: group:cn=Anybody:normal:rsc
aclSource: ou=deptXYZ, o=IBM, c=US
```

In the absence of an explicit ACL for entry cn=personA, ou=deptXYZ, o=IBM, c=US, the following is the implicit, propagated ACL for the entry:

```
aclPropagate: TRUE
aclEntry: group:cn=deptXYZRegs, o=IBM, c=US:normal:rcs:sensitive:rsc
aclEntry: access-id:cn=personA, ou=deptXYZ, o=IBM, c=US:object:ad:normal:rwsc:sensitive:rwsc:critical:rsc
aclEntry:
           group:cn=Anybody:normal:rsc
aclSource: ou=deptXYZ, o=IBM, c=US
```

In this example, a propagating ACL has been set on ou=deptXYZ, o=IBM, c=US. No ACL has been set on the descendant cn=personA, ou=deptXYZ, o=IBM, c=US. Therefore, the descendant inherits its ACL value from the nearest ancestor with a propagating ACL. This happens to be ou=deptXYZ, o=IBM, c=US, which is reflected in the aclSource attribute value. The aclEntry and aclPropagate values are identical to the explicit propagating ACL set at ou=deptXYZ, o=IBM, c=US.

Examples of overrides

Following is an explicit ACL for entry o=IBM, c=US:

```
aclPropagate: TRUE
aclEntry: group:cn=IBMRegs, o=IBM, c=US:normal:rcs:sensitive:rsc
aclEntry:
           group:cn=Anybody:normal:rsc
aclSource: o=IBM, c=US
```

Following is an explicit ACL for entry ou=deptXYZ, o=IBM, c=US:

```
aclPropagate: FALSE
aclEntry: group:cn=deptXYZRegs, o=IBM, c=US:normal:rcs:sensitive:rsc
         access-id:cn=personA, ou=deptXYZ, o=IBM, c=US:object:ad:normal:rwsc:sensitive:rwsc:critical:rsc
aclEntry:
aclEntry: group:cn=Anybody:normal:rsc
aclSource: ou=deptXYZ, o=IBM, c=US
```

Note that in the explicit ACLs above, aclSource is the same as the entry DN. This attribute is generated and managed by the directory server; it cannot be set when modifying ACLs.

Following is an implicit ACL for entry cn=personA, ou=deptXYZ, o=IBM, c=US:

```
aclPropagate: TRUE
aclEntry: group:cn=IBMRegs, o=IBM, c=US:normal:rcs:sensitive:rsc
aclEntry: group:cn=Anybody:normal:rsc
aclSource: o=IBM, c=US
```

In this example, a propagating ACL has been set on o=IBM, c=US. An override ACL has been set (acIPropagate is FALSE) on the descendant ou=deptXYZ, o=IBM, c=US. Therefore, the ACL set at ou=deptXYZ, o=IBM, c=US pertains only to that particular entry.

The descendant cn=personA, ou=deptXYZ, o=IBM, c=US inherits its ACL value from the nearest ancestor with a propagating ACL (which is o=IBM, c=US as reflected in the aclSource). The ACL on ou=deptXYZ, o=IBM, c=US is not used because aclPropagate is FALSE.

Other examples

In these examples, the administrator DN will be cn=admin, c=US.

The following example shows the default ACL:

```
aclPropagate: TRUE
aclEntry: group:cn=Anybody:normal:rsc:system:rsc
aclSource: default
ownerPropagate: TRUE
entryOwner: access-id:cn=admin,c=US
ownerSource: default
The following example shows a typical ACL for entry cn=personA, ou=deptXYZ, o=IBM, c=US:
aclPropagate: TRUE
aclEntry: group:cn=deptXYZRegs, o=IBM, c=US:normal:rcs:sensitive:rsc
aclEntry: access-id:cn=personA, ou=deptXYZ, o=IBM, c=US:object:ad:normal:rwsc
aclEntry: group:cn=Anybody:normal:rsc:system:rsc
aclSource: ou=deptXYZ, o=IBM, c=US
ownerPropagate: TRUE
entryOwner: access-id:cn=deptXYZMgr, ou=deptXYZ, o=IBM, c=US
ownerSource: ou=deptXYZ, o=IBM, c=US
```

This is an inherited ACL and an inherited owner. Both owner properties and ACL properties are inherited from entry ou=deptXYZ, o=IBM, c=US. In this example, members of group cn=deptXYZRegs, o=IBM, c=US have permission to read, search and compare entries in both the normal and sensitive attribute access classes. They do not have permission to add or delete entries under this entry. Nor do they have permission to access any information or change any information on attributes in the critical attribute access class. Unauthenticated, as well as all other bound users, have permission to read, search, and compare attributes in the **normal** and **system** attribute access classes only. The personA has add and delete permission on the entry; read, write, search, and compare permissions on normal and sensitive attributes; and read, search, and compare permission on critical attributes. The deptXYZMgr had full access to the entry since it is the owner of the entry. As always, the administrator also has unrestricted access to the entry.

Access control groups

Access control groups provide a mechanism for applying the same aclEntry attribute values to an entry for multiple users without having to create an explicit aclEntry for each user. This is done by defining an accessGroup entry in the directory server and by including users as members of an access control group and then defining an **aclEntry** value for the group.

Access control groups have an object class of accessGroup. The accessGroup is a subclass of groupOfNames.

Each group entry contains a multi-valued attribute consisting of member DNs. Groups cannot contain other group DNs.

Retrieving ACL information from the LDAP server

In order to retrieve all of the ACL information in a namespace, use the **Idapsearch** command, as shown in the following example:

```
Idapsearch -h 127.0.0.1 -D "cn=admin, dc=Your Company,dc=com" -w xxxxxx -b "dc=Your Company,dc=com" \
"(objectclass=*)" aclEntry aclPropagate aclSource entryOwner ownerPropagate ownerSource
dn:Your Company, dc=com
aclPropagate: TRUE
aclEntry: CN=ADMIN:normal:rwsc:sensitive:rwsc:critical:rwsc:object:ad
aclEntry: CN=ANYBODY:normal:rsc:system:rsc
aclSource: dc=Your Company, dc=com
ownerPropagate: TRUE
entryOwner: CN=ADMIN
ownerSource: default
```

This command performs a subtree search starting at the root of the tree (assuming the root of the tree is "dc=Your Company,c=com") and returns the six ACL attributes for each entry in the tree. It is necessary to

specifically request the six ACL attributes because they are considered as "operational" and, therefore, can only be returned on a search if requested. (See IETF RFC 2251, The Lightweight Directory Access Protocol (V3).)

ACL information (aclEntry, aclPropagate, aclSource, entryOwner, ownerPropagate, and ownerSource) is returned for all entries. For those entries that contain ACLs, the aclSource and ownerSource attributes contain the same DN as the entry DN. For those entries that do not contain ACLs, the aclSource and ownerSource attributes contain distinguished names of the entries that contain the ACL information (aclEntry and entryOwner) that are used for access control checking of information in that entry.

Notes:

- 1. It is possible for the aclSource and ownerSource attributes to contain the value default. This is not a distinguished name but rather represents that the ACL that applies to the entry is the default ACL (that is, no ACL information has been specified to apply to the entry).
- 2. If the tree is larger than the **sizeLimit** setting in the LDAP server configuration file or on the serach command, then not all entries are returned. In this case, use adminDN to look for the information.

You can also use the same method to get the ACL information for a portion of the namespace by specifying the -b searchbase parameter on the Idapsearch command, where searchbase is the starting point for the search.

Creating and managing access controls

To create and update ACLs in the TDBM backend, use a tool implementing Idap_modify APIs, such as Idapmodify. The Idapmodify program is a general directory tool that allows creation, modification, and deletion of any set of attributes that are associated with an entry in the directory. Since access control information is maintained as a set of additional attributes within an entry, Idapmodify is a natural choice for administering access control information in the TDBM database backend.

See z/OS: Security Server LDAP Client Programming for details on using the LDAP utilities, such as Idapmodify.

Creating an ACL

In order to create an ACL, the aclEntry and aclPropagate attributes must be added to the information stored for an entry. The aclEntry and aclPropagate attributes are added to an entry by either specifying them as part of the entry information when the entry is added to the directory or by modifying the entry after it exists to contain the aclEntry and aclPropagate information.

It is possible to create an ACL without specifying the aclPropagate attribute. In this case, the aclPropagate attribute is assumed to have a value of TRUE and is added into the directory entry automatically, along with the aclEntry information.

Since the **Idapmodify** command is very powerful, all the possible ways of adding the **aclEntry** and aclPropagate information cannot be shown here. The examples shown here describe the more common uses of the **Idapmodify** command to add ACL information.

Figure 37 on page 260 shows how to add a propagating ACL with three aclEntry values to an existing entry replacing any current aclEntry value .

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f newAcl.ldif
Where newAcl.ldif contains:
dn: cn=tim, o=Your Company
changetype: modify
replace: aclentry
aclEntry: cn=jeanne, o=Your Company:
normal:rsc:sensitive:rsc:critical:rsc
aclEntry: cn=jeff, cn=tim, o=Your Company:normal:rsc
aclEntry: cn=tim, o=Your Company:
normal:rwsc:sensitive:rwsc:critical:rwsc
```

Figure 37. Example of adding propagating ACL to existing entry in directory

The ACL added in Figure 37 is created as a propagating ACL since the aclPropagate attribute is not specified and so assumed to be TRUE. This means that the ACL will apply to all entries below cn=tim, o=Your Company that do not already have an ACL associated with them. Note that the first and last aclEntry values span two lines in the newAcl.ldif file. In order to do this, the first character on the continued line must be a space character, as shown in the example.

While it is not required that the LDAP administrator update all ACL information, the examples in this section all use the administrator when updating ACLs. Further, the use of -h 127.0.0.1 implies that the Idapmodify commands are performed from the same system on which the LDAP server is running and that the LDAP server is listening on TCP/IP port 389. Refer to the Idapmodify command description in z/OS: Security Server LDAP Client Programming for more details on the -h, -D, and -w command-line options. The ACL attributes can be updated from any LDAP client as long as the user performing the updates has the proper authorization to update the ACL information.

The ACL attributes are defined to be in a special access class called restricted. Thus, in order to allow someone other than the LDAP administrator to update the ACL attributes, they must either be the entry owner or have the proper authorization to restricted attributes. Figure 38 shows an example of adding an ACL so that cn=jeanne, o=Your Company can update the ACL information:

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f newAcl.ldif
Where newAcl.ldif contains:
dn: cn=jeanne, o=Your Company
changetype: modify
replace: aclentry
aclEntry: cn=jeanne, o=Your Company:
normal:rsc:sensitive:rsc:critical:rsc:restricted:rwsc
aclEntry: cn=jeff, cn=tim, o=Your Company:normal:rsc
aclEntry: cn=tim, o=Your Company:
normal:rsc
```

Figure 38. Example of adding propagating ACL to existing entry in the directory. In this example, one aclEntry value allows a person update capability on ACL information.

The ACL added in Figure 38 allows cn=jeanne, o=Your Company to update this ACL information. In addition, since the ACL is a propagating ACL, this allows cn=jeanne, o=Your Company to create new ACL information against any entry that is controlled by this ACL. Care must be taken here, however, since it is possible for cn=jeanne, o=Your Company to set up an ACL which then does not allow cn=jeanne, o=Your Company update capability on the ACL information. If this occurs, a user with sufficient authority (the administrator, for example) must be used in order to reset/change the ACL information.

Figure 39 shows an example of adding a non-propagating ACL. A non-propagating ACL applies only to the entry to which it is attached and not to the subtree of information that might be stored below the entry in the directory.

```
$ ldapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f newAcl.ldif
Where newAcl.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
replace: aclentry
aclEntry: cn=tim, o=Your Company:normal:rwsc:sensitive:rwsc:
critical:rwsc:restricted:rwsc
aclEntry: cn=jeff, cn=tim, o=Your Company:normal:rwsc:
sensitive:rwsc:critical:rwsc
aclEntry: cn=jeanne, o=Your Company:normal:rsc
replace: aclpropagate
aclPropagate:FALSE
```

Figure 39. Example of setting up a non-propagating ACL

Setting up a non-propagating ACL is similar to setting up a propagating ACL. The difference is that the aclPropagate attribute value is set to FALSE.

Modifying an ACL

Once an ACL exists for an entry in the directory, it may have to be updated. To do this, the **Idapmodify** command is used. As described earlier in this chapter, while the Idapmodify command is used in these examples, what is really being used is an LDAP client application, issuing LDAP modify operations to the LDAP server. Thus, modifications to ACL information need not be performed from the same system on which the LDAP server is running.

Modifications to ACLs can be of a number of different types. The most common modifications are to:

- Add an additional aclEntry value to the ACL to allow another person or group access to the entry
- Change an ACL from propagating to non-propagating
- Remove an aclEntry value which exists in the ACL to disallow another person or group access to the entry that they had before.

Figure 40, Figure 41 on page 262, and Figure 42 on page 262 show examples of these modifications, respectively.

Figure 40 shows how an additional **aclEntry** value is added to existing ACL information.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modAcl.ldif
Where modAcl.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
add: aclentry
aclEntry: cn=dylan, cn=tim, o=Your Company:
normal:rwsc:sensitive:rwsc:critical:rwsc:restricted:rwsc
```

Figure 40. Example of adding an aclEntry attribute value

In Figure 40 on page 261, cn=dylan, cn=tim, o=Your Company is granted permissions against the cn=jeff, cn=tim, o=Your Company entry in the directory. The existing ACL information remains in the entry; the aclEntry attribute value for cn=dylan, cn=tim, o=Your Company is added to this information.

Figure 41 shows how to modify an existing ACL to be non-propagating instead of propagating.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modAcl.ldif
Where modAcl.ldif contains:
dn: cn=tim, o=Your Company
changetype: modify
replace: aclpropagate
aclPropagate: FALSE
```

Figure 41. Example of modifying aclPropagate attribute

In Figure 41, the existing ACL against cn=tim, o=Your Company is modified to be a non-propagating ACL instead of a propagating ACL. This means that the ACL will no longer apply to entries below cn=tim, o=Your Company in the directory tree. Instead, the first propagating ACL that is found in an entry above cn=tim, o=Your Company will be applied to the entries below cn=tim, o=Your Company. If no ACL is found in the entries above cn=tim, o=Your Company, then the default ACL is used.

Figure 42 shows how to remove an **aclEntry** value from existing ACL information:

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modAcl.ldif
Where modAcl.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
delete: aclentry
aclEntry: cn=dylan, cn=tim, o=Your Company
```

Figure 42. Example of removing a single aclEntry attribute value

In Figure 42, the aclEntry attribute value for cn=dylan, cn=tim, o=Your Company is removed from the ACL information for entry cn=jeff, cn=tim, o=Your Company.

Deleting an ACL

In order to delete an ACL that is attached to an entry in the directory, the aclEntry and aclPropagate attributes must be deleted from the entry. To do this, use the Idapmodify command to delete the entire attribute (all values) from the entry.

Figure 43 on page 263 shows an example of deleting an ACL from an entry.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f delAcl.ldif
Where delAcl.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
delete: aclentry
delete: aclPropagate
```

Figure 43. Example of deleting an ACL from an entry

In Figure 43, the existing ACL against cn=jeff, cn=tim, o=Your Company is removed. This means that the ACL will no longer apply to the entry. Instead, the first propagating ACL that is found in an entry above cn=jeff, cn=tim, o=Your Company will be applied to cn=jeff, cn=tim, o=Your Company. If no ACL is found in the entries above cn=jeff, cn=tim, o=Your Company, then the default ACL is used.

Note: Both the aclEntry and aclPropagate attributes must be deleted from the entry even though the aclPropagate attribute might not have been specified during the creation of the ACL information.

Creating an owner for an entry

In addition to the access control list control of directory entries, each entry can have assigned to it an entry owner or set of entry owners. As an entry owner, full access is allowed to the entry. Entry owners are granted add and delete permission, as well as read, write, search, and compare for all attribute classes. Entry owners can add and modify ACL information on the entries for which they are specified as the owner.

Entry owners are listed in the entryOwner attribute. Just like aclEntry information, entryOwner information can be propagating or non-propagating based on the setting of the ownerPropagate attribute. Like the aclSource attribute for aclEntry information, the ownerSource attribute lists the distinguished name of the entry that contains the entryOwner attribute which applies to the entry.

In order to create an entry owner, the entry Owner and owner Propagate attributes must be added to the information stored for an entry. The entryOwner and ownerPropagate attributes are added to an entry by either specifying them as part of the entry information when the entry is added to the directory or by modifying the entry after it exists to contain the entryOwner and ownerPropagate information.

It is possible to create an entry owner without specifying the **ownerPropagate** attribute. In this case, the ownerPropagate attribute is assumed to have a value of TRUE and is added into the directory entry automatically.

Since the Idapmodify command is very powerful, all the possible ways of adding the entryOwner and ownerPropagate information cannot be shown here. The examples shown here describe the more common uses of the **Idapmodify** command to add entry owner information.

Figure 44 on page 264 shows how to add a propagating entry owner with two entryOwner values to an existing entry.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f newOwn.ldif
Where new0wn.ldif contains:
dn: cn=tim, o=Your Company
changetype: modify
replace: entryOwner
entryOwner: cn=joe, o=Your Company
entryOwner: cn=carol, o=Your Company
```

Figure 44. Example of adding a propagating set of entry owners to existing entry in the directory

The entry owners added in Figure 44 are created as a propagating set of entry owners since the ownerPropagate attribute is not specified and so assumed to be TRUE. This means that the entry owners will apply to all entries below cn=tim, o=Your Company that do not already have an entry owner associated with them.

While it is not required that the LDAP administrator update all entry owner information, the examples in this section all use the administrator as the entry owner updater. Further, the use of -h 127.0.0.1 implies that the Idapmodify commands are performed from the same system on which the LDAP server is running and that the LDAP server is listening on TCP/IP port 389. Refer to the Idapmodify command description in z/OS: Security Server LDAP Client Programming for more details on the -h, -D, and -w command-line options. The entry owner attributes can be updated from any LDAP client as long as the user performing the update has the proper authorization to update the entry owner information.

The entry owner attributes, like the ACL attributes, are defined to be in a special access class called restricted. Thus, in order to allow someone other than the LDAP administrator to update the entry owner attributes, they must have the proper authorization to restricted attributes. See Figure 38 on page 260 for an example of allowing users other than the LDAP administrator the ability to update entry owner information.

Figure 45 shows an example of adding a non-propagating entry owner. A non-propagating entry owner applies only to the entry to which it is attached and not to the subtree of information that might be stored below the entry in the directory.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f newOwn.ldif
Where new0wn.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
replace: entryOwner
entryOwner: cn=george, o=Your Company
entryOwner: cn=jane, o=Your Company
replace: ownerPropagate
ownerPropagate: FALSE
```

Figure 45. Example of setting up a non-propagating entry owner

Setting up a non-propagating entry owner is similar to setting up a propagating entry owner. The difference is that the **ownerPropagate** attribute value is set to **FALSE**.

Modifying an owner for an entry

Once an entry owner exists for an entry in the directory, it may have to be updated. To do this, the Idapmodify command is used. As described earlier in this chapter, while the Idapmodify command is used in these examples, what is really being used is an LDAP client application, issuing LDAP modify

operations to the LDAP server. Thus, modifications to entry owner information need not be performed from the same system on which the LDAP server is running.

Modifications to entry owners can be of a number of different types. The most common modifications are to:

- · Add an additional entryOwner value to the set of entry owners to allow another person or group to control the entry
- Change an entry owner from propagating to non-propagating
- Remove an entryOwner value which exists in the entry owner set to disallow another person or group access to control the entry that they had control over before.

Figure 46, Figure 47, and Figure 48 on page 266 show examples of these modifications, respectively.

Figure 46 shows how an additional **entryOwner** value is added to existing entry owner information.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modOwn.ldif
Where mod0wn.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
add: entryOwner
entryOwner: cn=george, o=Your Company
```

Figure 46. Example of adding an entryOwner attribute value

In Figure 46, cn=george, o=Your Company is granted entry owner control of the cn=jeff, cn=tim, o=Your Company entry in the directory. The existing entry owner information remains in the entry; the entryOwner attribute value for cn=george, o=Your Company is added to this information.

Figure 47 shows how to modify an existing entry owner to be non-propagating instead of propagating.

```
$ ldapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modOwn.ldif
Where mod0wn.ldif contains:
dn: cn=tim. o=Your Company
changetype: modify
replace: ownerPropagate
ownerPropagate: FALSE
```

Figure 47. Example of modifying the ownerPropagate attribute

In Figure 47, the existing entry owner set for cn=tim, o=Your Company is modified to be a non-propagating instead of a propagating. This means that the entry owner will no longer apply to entries below cn=tim, o=Your Company in the directory tree. Instead, the first propagating entry owner set that is found in an entry above cn=tim, o=Your Company will be applied to the entries below cn=tim, o=Your Company. If no entry owner is found in the entries above cn=tim, o=Your Company, then the default entry owner is used.

Figure 48 on page 266 shows how to remove an **entryOwner** value from existing entry owner information:

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modOwn.ldif
Where mod0wn.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
delete: entryOwner
entryOwner: cn=george, cn=tim, o=Your Company
```

Figure 48. Example of removing a single entryOwner Attribute value

In Figure 48, the **entryOwner** attribute value for cn=george, cn=tim, o=Your Company is removed from the entry owner information for entry cn=jeff, cn=tim, o=Your Company.

Deleting an owner for an entry

In order to delete an entry owner set that is attached to an entry in the directory, the entryOwner and ownerPropagate attributes must be deleted from the entry. To do this, use the Idapmodify command to delete the entire attribute (all values) from the entry.

Figure 49 shows an example of deleting an entry owner set from an entry.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f delOwn.ldif
Where delown.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
delete: entryOwner
delete: ownerPropagate
```

Figure 49. Example of deleting an entry owner set from an entry

In Figure 49, the existing entry owner set against cn=jeff, cn=tim, o=Your Company is removed. This means that the entry owner information will no longer apply to the entry. Instead, the first propagating entry owner set that is found in an entry above cn=jeff, cn=tim, o=Your Company will be applied to cn=jeff, cn=tim, o=Your Company. If no entry owner set is found in the entries above cn=jeff, cn=tim, o=Your Company, then the default entry owner is used.

Note: The entryOwner and ownerPropagate attributes must be deleted from the entry even though the ownerPropagate attribute might not have been specified during the creation of the entry owner information.

Creating a group for use in ACLs and entry owner settings

Sets of users can be grouped together in the directory by defining them as members of a group in the directory. A directory group, used for access control checking, is just another entry in the directory that is defined to have an object class value of accessGroup. Access controls can be defined using groups of users.

Note: Deleting a user or a group does not have any cascade effect on any ACLs that include that user or group. The user or group DN is not removed from the ACLs. If another user or group is subsequently created with the same DN, that user or group will be granted the privileges of the former user or group.

An accessGroup entry contains a multi-valued attribute called member. The member attribute contains the distinguished names of all the members of the group. A user is said to be "in a group" if the member attribute for the group contains a value that is the distinguished name of the entry that represents the user.

When defining access controls or entry owner sets, names of accessGroup entries can be used in the same place as user entry names. When access control decisions are performed, a user's accessGroup memberships are used in determining if a user can perform the action requested.

Figure 50 shows an example of creating an accessGroup in the directory. The accessGroup entry contains four users in the group. The name of the accessGroup can then be used in aclEntry and **entryOwner** attributes when setting up access control information.

```
$ ldapmodify -a -h 127.0.0.1 -D "cn=admin" -w xxxx -f newGrp.ldif
Where newGrp.ldif contains:
dn: cn=group1, o=Your Company
objectclass: accessGroup
member: cn=bob, o=Your Company
member: cn=lisa, o=Your Company
member: cn=chris, cn=bob, o=Your Company
member: cn=john, cn=bob, o=Your Company
```

Figure 50. Example of creating an accessGroup entry in the directory

Note that while the accessGroup entry will allow any distinguished name to be added as a member attribute value, the access control implementation in the TDBM backend does not support groups to contain other groups as members.

Groups are added to access control information in just the same way as user entries are added to access control information. Figure 51 shows how the group that was defined in Figure 50 can be added to the aclEntry information in an existing access control specification for an entry. Figure 52 shows how the group defined in Figure 50 can be added as an **entryOwner** to an existing entry owner specification for an entry.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modAcl.ldif
Where modAcl.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
add: aclentry
aclEntry: cn=group1, o=Your Company normal:rwsc:sensitive:rsc
Figure 51. Example of adding a group to access control information
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modOwn.ldif
Where mod0wn.ldif contains:
dn: cn=jeff, cn=tim, o=Your Company
changetype: modify
add: entryOwner
entryOwner: cn=group1, o=Your Company
```

Figure 52. Example of adding a group to entry owner information

Modifying a group

In order to modify an accessGroup, use the Idapmodify command. Users can be added or removed from the group individually. Figure 53 and Figure 54 show examples of adding and removing members from an accessGroup entry, respectively.

In order to add a member to a group, add the user's distinguished name as an additional value for the member attribute. Figure 53 shows an example for doing this using the **Idapmodify** command.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modGrp.ldif
Where modGrp.ldif contains:
dn: cn=group1, o=Your Company
changetype: modify
add: member
member: cn=jeff, cn=tim, o=Your Company
```

Figure 53. Adding a member to an accessGroup

In order to remove a member from a group, remove the user's distinguished name from the set of member attribute values in the accessGroup entry. Figure 54 shows an example of this using the Idapmodify command.

```
$ Idapmodify -h 127.0.0.1 -D "cn=admin" -w xxxx -f modGrp.ldif
Where modGrp.ldif contains:
dn: cn=group1, o=Your Company
changetype: modify
delete: member
member: cn=jeff, cn=tim, o=Your Company
```

Figure 54. Deleting a member from an accessGroup

Deleting a group

In order to delete a group entry, delete the directory entry that represents the group. The Idapdelete command can be used to perform this delete operation. As an alternative, an LDAP client that supports the LDAP delete operation can be used to delete the entry from the directory.

An example of deleting an accessGroup entry from the directory is shown in Figure 55.

```
$ | ldapdelete -h 127.0.0.1 -D "cn=admin" -w xxxx "cn=group1, o=Your Company"
```

Figure 55. Deleting an entire accessGroup entry from the directory

Chapter 24. Replication

Once the z/OS LDAP server is installed and configured, users can access the directory, add objects, delete objects, or perform search operations to retrieve particular sets of information.

Replication is a process which keeps multiple databases in sync. Through replication, a change made to one database is propagated to one or more additional databases. In effect, a change to one database shows up on multiple different databases. This means there are two types of databases: masters and replicas.

Master

All changes to the database are made to the master server. The master server is then responsible for propagating the changes to all other databases. It is important to note that while there can be multiple databases representing the same information, only one of those databases can be the master.

Replica

Each of the additional servers which contain a database replica. These replica databases are identical to the master database.

Replication from a Master server to a Replica is only supported when the Master server is running in single-server mode. Although replication is not supported when operating multiple concurrent server instances against the same database (multi-server operating mode), similar benefits are afforded when operating in this mode. Refer to "Determining operational mode" on page 75 for more information about server operating modes.

In z/OS LDAP, replication is only supported in the TDBM (DB2-based) backend.

Ibm-entryuuid replication

Replication of the **ibm-entryuuid** will be performed to:

- OS/390 V2R10 and z/OS V1R1 with the Rollback PTF OW50971 and Coexistence PTF OW54426 applied
- z/OS V1R2 and z/OS V1R3 with the Coexistence PTF applied
- z/OS V1R4.

If a replica server is not one of the above, then the **ibm-entryuuid** attribute will not be replicated to that server; however, the entry and other attributes will be replicated.

Complex modify DN replication

Replication of Modify DN new superior operations will be performed to z/OS V1R4 LDAP replica servers. If a replica server is not at a supported level, the Modify DN new superior operation and all subsequent replication operations will be held until the replica is removed from the replica collection or it is upgraded to a LDAP z/OS V1R4 server.

Password encryption and replication

To ensure data integrity and the proper working of the LDAP servers in the replication environment, the **pwEncryption** option in the configuration files for the master and replica servers must be the same. If one of the servers involved in replication is a non-z/OS server, then the administrator must choose a **pwEncryption** method that is supported by both servers for correct operation of replication. If no encryption methods are common between the servers, then password encryption should not be used.

For crypt encryption, note that the values returned by the crypt algorithm are not portable to other X/Open-conformant systems. This means the crypt() algorithm cannot be used for replication between z/OS and a non-z/OS server.

For DES encryption, where both the master and replica servers are z/OS LDAP servers, the same DES key label and data key must be defined on both z/OS systems through the ICSF KGUP and CKDS facilities. (See the information on managing cryptographic keys in the z/OS: ICSF Administrator's Guide for more details.) This key label must be used in the configuration files of both of the LDAP servers involved in replication.

Benefits of replication

There are several benefits realized through replication. The single greatest benefit is providing a means of faster searches. Instead of having all search requests directed at a single server, the search requests can be spread among several different servers. This improves the response time for the request completion.

Additionally, the replica provides a backup to the master server. Even if the master server crashes, or is unreadable, the replica can still fulfill search requests, and provide access to the data.

Master server

In order for the replication process to occur, the following must happen:

- The master must be aware of each replica that is to receive the change information.
- Each replica must be aware of the master server for the database that it serves.

The master server becomes aware of the existence of the replica databases when objects of type replicaObject are added to the directory. Each of these objects represents a particular replica server. The attribute/value pairs within the replica object provide the information the server needs in order to find the replica and send any updates to that server.

Replica objects

The replicaObject object class is provided in the system schema file schema.user.ldif. Like other LDAP object class definitions, the replicaObject has mandatory and optional attributes. Each of the replicaObject attributes are single-valued. The following is a description of the mandatory attributes of replicaObject.

Table 45. Replica object-schema definition (mandatory attributes)

Attribute	Description and example
replicaHost	Represents the Internet name of the machine on which the replica resides. This could be an IP address, such as, 127.0.0.1, or a DNS name such as myMachine.endicott.ibm.com.
	Example:
	replicaHost: myMachine.endicott.ibm.com
replicaBindDN	Specifies the LDAP distinguished name that the master uses to bind to the replica when sending directory updates. The replicaBindDN and the masterServerDN in the replica's configuration file must have the same value.
	Example:
	replicaBindDN: cn=Master
replicaCredentials	Contains the authentication information needed for the master server to authenticate to the replica using the replicaBindDN .
	Example:
	replicaCredentials: secret

Table 45. Replica object-schema definition (mandatory attributes) (continued)

Attribute	Description and example	
cn	Forms the RDN of the LDAP distinguished name of the replicaObject entry.	
	Example:	
	cn: myReplica	

In the examples in Table 45 on page 270, when the master server receives and successfully finishes an update request, the update is also sent to myMachine.endicott.ibm.com on port 389 (the default port). The master performs a bind operation using the DN of cn=Master and password of secret.

In addition, there are several attributes available that provide additional flexibility in configuring a replica server. For instance, an added description could better describe the replica server, and it could listen on a different port then the default port of 389. Examples of adding a description and changing the port to 400 are shown in Table 46, which describes the optional attributes of replicaObject.

Table 46. Replica object schema definition (optional attributes)

Attribute	Description and example
replicaPort	Describes the port number on which the replica is listening for incoming requests. By default, the server listens on port 389.
	Example:
	replicaPort: 400
replicaUpdateTimeInterval	Delays the propagation of additional updates for specified number of seconds. The default is for the master server to send updates immediately.
	Example:
	replicaUpdateTimeInterval: 3600
replicaUseSSL	Determines whether the master should replicate over SSL/TLS. The default is for the master not to replicate using SSL/TLS.
	Example:
	replicaUseSSL: TRUE
description	Provides an additional text field for extra information pertaining to the replica object.
	Example:
	description: Replica Machine in the fourth floor lab
seeAlso	Identifies another directory server entry that may contain information related to this entry.
	Example:
	seeAlso: cn=Alternate Code, ou=Software, o=IBM, c=US
replicaBindMethod	Identifies the bind method to be used. If it is specified, it must be set to simple.
	Example:
	replicaBindMethod: simple

Replication only supports simple authentication. SASL EXTERNAL, GSSAPI, DIGEST-MD5, and CRAM-MD5 bind mechanisms are not supported as valid replication bind mechanisms.

Localhost suffix

When the LDAP server is configured with TDBM, the cn=localhost suffix is no longer required. However, its use is still supported. The **cn=localhost** suffix entry will not be created automatically in TDBM.

Adding replica objects in TDBM

In TDBM, replica objects can be placed anywhere within the directory tree. This also implies that the suffix cn=localhost can be removed from the LDAP server configuration file. Placing replica objects in the directory tree then requires that any parent entries of the replicaObject entry be added to the directory prior to adding the replicaObject entry. These entries must be added to both the master server and replica server before addition of the replicaObject. This is needed on the replica server because these entries are being added at the master server without replication being active. If a replica object is not placed as a leaf node in the directory tree, the only entries allowed below the replica object are other replica objects. The LDAP server will allow non-replica entries to be placed below replica entries; however, these entries will not be replicated to the slave. Following is an example of a replica object definition using LDIF format.

```
dn: cn=myReplica,o=YourCompany
objectclass: replicaObject
cn: myReplica
replicaHost: myMachine.endicott.ibm.com
replicaBindDn: cn=Master
replicaCredentials: secret
replicaPort: 400
replicaUseSSL: FALSE
description: "Replica machine in the fourth floor lab."
```

Replica server

Initialization, or population, of a replica database requires several steps.

Special note: If the master server is configured with TDBM, changes to the schema entry on the master are not replicated. The schema on the replica must be modified by a user bound as the masterServerDN. See "Configuring the replica" on page 273 for more information. A separate update of the replica schema will be required each time the schema is updated on the master. Note that if you are modifying the schema on a TDBM replica and are not bound as the masterServerDN, the masterServer configuration option will cause the modification to be redirected to the master. This will cause the replica and master schema to be out of sync. No error message occurs.

Populating a replica

- 1. Stop the LDAP master server.
- 2. Unload the master server's directory contents if there are any entries. For TDBM, use the tdbm2ldif utility (see "tdbm2ldif program" on page 137).
- 3. Make sure the schema for the replica server is the same as the schema for the master server. If the replica and master are both z/OS servers configured with TDBM, the schema can be unloaded from the master using tdbm2ldif and reloaded into the replica using either the ldif2tdbm -s option or **Idapmodify** with the replica server started.
- 4. Run a load utility with a single added directory entry which defines a replicaObject entry into the master server's directory contents. For TDBM, use either the Idif2tdbm utility (see "Idif2tdbm program" on page 127) or **Idapadd** with the master server running.
 - Note that in order to load the replicaObject entry, it is also necessary to load any parent entries in the directory hierarchy in hierarchy order.
- 5. If the master database does not contain any entries, no further action must be taken to ensure that the replica and master server are in sync and the master server can now be restarted; otherwise, continue to the next step.

- 6. Transport the LDIF file created in step 2 on page 272 to the replica server's location.
- 7. Run a load utility on the replica server using the LDIF file from step 6. For TDBM, stop the replica server if it is running and use Idif2tdbm.
- 8. Configure the replica (see next section).
- 9. Start the replica server.
- 10. Start the master server.

Configuring the replica

The key to a successful replica configuration rests in ensuring that the values in the replicaObject on the master server accurately represent the relevant values on the replica server. Configuring the replica involves specifying appropriate configuration file option values to identify:

- the IP address and port on which the replica server should listen for communication from the master server
- the type of connection expected by the master server when it communicates to the replica server, either over a non-secure or secure connection
- the DN and password used by the master server.

Note: I Idif2tdbm does not replicate changes to a replica. So, if you are using Idif2tdbm to add entries to a masterpiece you must also use it to add entries to the replica, with no intervening updates on the master before the replica is loaded.

The following table identifies the relationship between replicaObject on a z/OS LDAP master server and the configuration options on an IBM replica server. The values specified for these options must be equivalent. An example of what is meant by equivalent is when the replica server is listening on all of its network interfaces, then the replicaHost must specify either the corresponding hostnname or an IP address of one of the addresses:

Attribute in replica object on master server	Corresponding replica server configuration option or command line parameter
replicaHost	listen configuration option or -I LDAP server command line parameter
replicaPort	listen , port , or securePort configuration options -p , -s , or -I LDAP server command line parameter
replicaUseSSL	listen configuration options -p , -s , or -l LDAP server command line parameter
replicaBindDn	masterServerDN
replicaCredentials	masterServerPw

Notes:

- 1. Listen configuration option is used only by the z/OS LDAP server.
- 2. If the replica server is a non-IBM server, you should consult their documentation for parameters that correspond to the parameters mentioned in the above table.
- 3. If a replica server is configured with more than one TDBM backend then the masterServer, masterServerDN, and masterServerpw all have to appear in each TDBM backend sections in the configuration file and all must have the same values. See "Establishing the administrator DN and password" on page 80 for details.
- 4. It is recommended that the masterServerDN be a DN that is dedicated specifically to replication. It should not be used for any other operations.
- 5. The masterServer, masterServerDN, and masterServerpw entries must follow the database definition entry in the LDAP server configuration file.

LDAP update operations on replicas

Update operations, such as add, delete, modify, and rename, should not be performed against a replica server. Changes must be made to the master database, which then propagates the change to the replica.

If update operations are sent to a replica server, the masterServer set in the replica configuration is returned and the operation is referred to the master server and is then propagated to the replica server. For compatibility purposes, if masterServer is not specified, the first default referral found in the replica configuration file is used as the master server and all update operations are returned to that server.

In order to maintain database integrity, a load utility (Idif2tdbm) should be used on a replica only when initially populating the replica database. If a load utility is used to add entries to a replica server after initial population, these changes are not reflected in the master database. The replica database is corrupted and could give erroneous information. See "SSL/TLS and replication" for information about securing a database.

Changing a replica to a master

At some point, it may become desirable to change one of the replicas to be the master. Perhaps the machine where the replica server is installed is being upgraded, and the customer wishes this replica to now be the master LDAP server.

The following procedure should be followed to change a replica to a master:

- 1. Ensure all of the data (including replicaObject entries for each replica object) from the master resides in the replica. This can be done by dumping both databases using an unload utility and comparing the output. (For TDBM, use tdbm2ldif.) If the replica is out of sync with the master, follow the procedure for correcting out-of-sync conditions.
- 2. Remove the **replicaObject** entry for this replica from the master server.
- 3. Stop the master and the replica servers.
- 4. Remove the masterServer, masterServerDN, and masterServerPw directives from the replica's configuration file.
- 5. If the original master is being eliminated, simply drop the databases on the original master. (See "Creating the DB2 database and table spaces for TDBM" on page 40 for examples of the SPUFI commands needed to drop the databases.) The new master can now be started.
- 6. If the original master is going to become a replica:
 - a. Add a replicaObject for the original master to the new master database using Idapadd or a load utility (Idif2tdbm for TDBM).
 - b. Add the masterServer directive to the new replica's configuration file. Make sure it points to the new master.
 - c. Add the masterServerDN and masterServerPw directives to the new replica's configuration file.
 - d. Start the servers.

SSL/TLS and replication

SSL/TLS can be used to communicate between a Master and Replica LDAP server.

Replica server with SSL/TLS enablement

Set the replica server up for SSL/TLS just like a normal SSL/TLS server. It needs its own public-private key pair and certificate, and the configuration file needs the standard SSL keywords set (sslKeyRingFile and sslKeyRingFilePW configuration file options). See "Setting up for SSL/TLS" on page 43 for more information.

Master server with SSL/TLS enablement

The master server acts like an SSL/TLS client to the replica server.

To set up the master server, you must:

- 1. Run the **gskkyman** utility (see *z/OS: System Secure Sockets Layer Programming*), this time as if you were a client. You should use the same key database file that contains the master server's key pair and certificate. Receive the replica's self-signed certificate and mark it as trusted.
- 2. In the master server's configuration file:
 - Set the sslKeyRingFile to the replica key database file name created above.
 - Set the sslKeyRingFilePW to the password for the replica key database file, or set sslKeyRingPWStashFile to the file name where the password is stashed.
- 3. In the replica object:
 - Set the replicaPort keyword to the replica's secure port number.
 - Set the replicaUseSSL keyword to TRUE.

See "Setting up for SSL/TLS" on page 43 for more information.

Since the master server acts like an SSL/TLS client to the replica server, the master server binds with the replica server. The bind method used is simple bind. The SASL external bind method is not supported for replication.

Troubleshooting

If the replica server does not seem to be receiving updates from the master server, there are several possible reasons. Check the following conditions for a possible guick fix:

- · Check for messages from the master server.
- · Verify that the replica object exists in the master database, and was specified correctly to match with the replica server. If cn=localhost is used as the base for all replica objects, perform an Idapsearch with a base of cn=localhost and a filter of objectClass=*.

If using TDBM and replica objects can be anywhere in the namespace, perform an Idapsearch where the search base is the suffix defined in the configuration file and the filter is objectClass=replicaObject. If more than one suffix is configured for TDBM, the search must be repeated using each suffix in the search base.

See z/OS: Security Server LDAP Client Programming for more information about Idapsearch.

- · Verify that the replicaHost value in the replica object for that replica specifies the machine on which the replica is running.
- · Check that the values listed in the replica object for that replica match those of the replica server configuration. Specifically, the replicaPort, replicaBindDN, and replicaCredentials should be verified.
- · Check that the replicaUpdateTimeInterval specified in the replica object for that replica has been set correctly.
- Verify that the replica server is running by performing an Idapsearch against the replica.
- · Check that the default referral specified in the replica's configuration file points to the master server.
- If the replicaObject attribute replicaUseSSL is set to TRUE, verify the replicaObject attribute replicaPort is set to the SSL port configured on the replica server, and verify the sslKeyRingFile, and sslKeyRingFilePW or sslKeyRingPWStashFile values in the master's configuration file are correct.

Recovering from out-of-sync conditions

If a replica becomes out-of-sync with its master for any reason, and normal replication processing is not correcting the situation, it may be necessary to reload the replica.

The following procedure should be followed to reload a replica:

- 1. Delete the replicaObject entry or entries on the master. It is desirable to search these entries first, obtaining LDIF format output for these entries. This LDIF can be used to reload the replicaObjects in step 7.
 - 2. Stop both the master and replica servers.
 - 3. Drop and recreate the table spaces on the replica server. (See "Creating the DB2 database and table spaces for TDBM" on page 40 or for an example of the SPUFI commands needed to drop and recreate the table spaces.)
 - 4. Run an unload utility on the master. Use tdbm2ldif for TDBM to unload the schema and the database.
 - 5. Run a load utility on the replica, using the data retrieved from the master, above. Use Idif2tdbm for TDBM.
 - 6. Start both servers.
- 7. Add the replicaObject entry or entries back to the master. The LDIF output captured from the search in step 1 can be used to add the entries back into the master.

Chapter 25. Referrals

Referrals provide a way for servers to refer clients to additional directory servers. With referrals you can:

- Distribute namespace information among multiple servers
- · Provide knowledge of where data resides within a set of interrelated servers
- Route client requests to the appropriate server

Following are some of the advantages of using referrals:

- · Distribute processing overhead, providing primitive load balancing
- Distribute administration of data along organizational boundaries
- Provide potential for widespread interconnection, beyond an organization's own boundaries.

This chapter describes how to use the **referral** object class and the **ref** attribute to construct entries in an LDAP directory server containing references to other LDAP directory servers. Also described in this chapter is how to associate multiple servers using referrals and an example of associating a set of servers through referrals and replication (see Chapter 24, "Replication" on page 269).

Using the referral object class and the ref attribute

The **referral** object class and the **ref** attribute are used to facilitate distributed name resolution or to search across multiple servers. The **ref** attribute appears in an entry named in the referencing server. The value of the **ref** attribute points to the corresponding entry maintained in the referenced server. While the distinguished name (DN) in a value of the **ref** attribute is typically that of an entry in a naming context below the naming context held by the referencing server, it is permitted to be the distinguished name of any entry. A multi-valued **ref** attribute may be used to indicate different locations for the same resource. If the **ref** attribute is multi-valued, all the DNs in the values of the **ref** attribute should have the same value.

Creating entries

Following is an example configuration that illustrates the use of the **ref** attribute.

Server A dn: o=ABC,c=US objectclass: referral objectclass: extensibleObject ref: ldap://hostB/o=ABC,c=US dn: o=XYZ,c=US objectclass: referral objectclass: extensibleObject ref: ldap://hostC/o=XYZ,c=US ref: ldap://hostD/o=XYZ,c=US

Server B	Server C
dn: o=ABC,c=US o: ABC other attributes	dn: o=XYZ,c=US o: XYZ other attributes

Server D
dn: o=XYZ,c=US
o: XYZ
other attributes

Figure 56. Example using ref attribute

In the example, Server A holds references to two entries: o=ABC,c=US and o=XYZ,c=US. For the o=ABC,c=US entry, Server A holds a reference to Server B and for the o=XYZ,c=US entry, Server A holds references to two equivalent servers, Server C and Server D.

The recommended setup of referrals is to structure the servers into a hierarchy based on the subtrees they manage. Then, provide "forward" referrals from servers that hold higher information and set the default referral to point back to its parent server.

Associating servers with referrals

In order to associate servers through referrals:

- · Use referral objects to point to other servers for subordinate references
- Define the default referral to point somewhere else, typically to the parent server

These steps are defined below.

Pointing to other servers

Use referral objects to point to the other servers for subordinate references. That is, portions of the namespace below this server which it does not service directly.

Referral objects, like other objects, go in the TDBM backend. Referral objects consist of:

dn Specifies the distinguished name. It is the portion of the namespace served by the referenced server.

objectclass

Specifies referral. For entries in TDBM, also include the object class extensibleObject.

ref Specifies the LDAP URL of the server. This URL should consist of the Idap:// or Idaps:// identifier, the hostname:port, and a DN. The DN requires a slash (/) before it to delimit it from the hostname:port, and should match the DN of the referral object. The ref attribute may be multi-valued, with each value specifying the LDAP URL of a different server. When multiple values are used, each LDAP URL should contain the same DN, and each server should hold equivalent information for the portion of the namespace represented by the DN.

Following is an example:

```
o=IBM.c=US
objectclass: referral
objectclass: extensibleObject
             ldap://Host1:389/o=IBM,c=US
             ldap://Host2:389/o=IBM,c=US
ref:
ref:
             ldap://Host3:1389/o=IBM,c=US
```

The server can have any number of referral objects within its database. However, the objects must essentially be descendents of its suffix.

Defining the default referral

Define the default referral to point to another server which services other portions of the namespace unknown to the referencing server. The default referral can be used to point to:

- The immediate parent of this server (in a hierarchy)
- A "more knowledgeable" server, such as the uppermost server in the hierarchy
- A "more knowledgeable" server which possibly serves a disjoint portion of the namespace.

The default referral goes in the configuration file and not the backend. The default referral is described in the configuration file with the referral keyword and an LDAP URL. Multiple default referrals may be specified. However, each one specified is considered equivalent; that is, each server referenced by a default referral should present the same view of the namespace to its clients.

Note: The default referral LDAP URL does not include the DN portion. It should just have the Idap:// identifier and the *hostname*:port. For example:

```
referral ldap://host3.ibm.com:999
```

See "Operating in multi-server mode without dynamic workload management enabled" on page 77 and "Operating in multi-server mode with dynamic workload management enabled" on page 77 for information about setting up multiple default referrals.

SSL/TLS note: A non-secure client referral to a secure port is not supported. Also, a secure client referral to a non-secure port is not supported.

Processing referrals

When clients request information from servers which do not hold the needed data, servers can pass back referral URLs which indicate one or more other servers to contact. The clients can then request the information from the referenced server. The z/OS client API, by default, chases referrals returned from servers. However, client applications can suppress referral chasing through the Idap_set_option() API. This option's scope is the LDAP handle, so a client could open multiple connections to one or more servers, some of which would chase referrals automatically, and some of which would not.

Servers present the referral URLs differently depending on the LDAP protocol version being used by the client. Referrals are presented to LDAP Version 2 clients in the error string, as the protocol does not provide a specific mechanism for indicating referrals. In LDAP Version 3, protocol elements are specifically defined to allow servers to present referral information to clients.

Using LDAP Version 2 referrals

LDAP Version 2 referrals are presented as part of the error string passed back to the client. Since clients do not generally examine the error string on results indicating LDAP_SUCCESS, some other error code is needed. Thus, on wholly successful operations, the server passes back a result of LDAP_PARTIAL_RESULTS instead of LDAP_SUCCESS to indicate the presence of referral information. On any result other than LDAP SUCCESS, referral information might also be present in the error string. Note that some servers might return a result of LDAP_PARTIAL_RESULTS with no referral information if the server does not have a default referral defined and the client makes a request for a portion of the namespace outside the server (and not below it in the hierarchy).

The referral information in the error string looks like this:

ldap://hostname:port/DN\n

where Referral: is followed by a new line character (\n) and ldap://hostname:port/DN\n is an LDAP URL followed by a new line character. The ellipses (...) indicate a list of multiple referrals; that is, more LDAP URLs followed by new line characters. Multiple referrals are only presented for partial search results when it is necessary to contact more than one additional server to complete the entire request. This would indicate that multiple referral objects were found in the referencing server that matched the search criteria. The client contacts every server presented in the list to continue the search request. For referral objects that have multi-valued ref attributes, the server sends only one of the LDAP URLs to a client using LDAP Version 2 protocol. This is because there is no provision for distinguishing between equivalent servers to contact (as indicated by multi-valued ref attributes) and multiple servers which must be contacted to complete a search request.

Limitations with LDAP Version 2 referrals

One limitation of LDAP Version 2 referrals is the lack of support for alternate referrals. As noted earlier, a referral object can have a multi-valued ref attribute which indicates different servers which hold equivalent information. As will be seen below, it is only possible for the LDAP Version 2 referral mechanism to contain a single list of referrals for a given request. Thus, one cannot tell if it is a list of servers, all of which must be contacted to complete a request, a list of equivalent servers of which only one should be contacted, or a combination of both.

A second limitation of referrals in LDAP Version 2 is that operations can sometimes be ambiguous in their intent regarding whether the operation was targeted for "real" objects in the namespace, as opposed to the referral objects themselves. For searches, referral objects are only presented as referrals, since the usual intent of a search is to look at the real objects in the namespace. Server administrators must therefore use other means to examine existing referral objects, such as examining the database, or reviewing tdbm2ldif output. For update operations, default referrals for upward references are presented as referrals, so that read-only replica servers can forward update operations to the master replica. However, subordinate references indicated by a referral object are not followed for update operations, rather they operate on the referral object itself. This is necessary to allow an administrator the ability to delete or modify existing referral objects. Erroneous changes caused by misdirected update operations are generally avoided through access protection and schema rules.

Using LDAP Version 3 referrals

In LDAP Version 3, referrals are defined as part of the protocol. The LDAP Version 2 limitations mentioned above are overcome by elements of the protocol and extensions to the protocol. There are two methods of passing back referral information in the LDAP Version 3 protocol: referrals and search continuation references.

A result code of LDAP REFERRAL is presented by the server to indicate that the contacted server does not hold the target entry of the request. The referral field is present in the result message and indicates another server (or set of servers) to contact. Referrals can be returned in response to any operation request except unbind and abandon which do not have responses. When multiple URLs are present in a given referral response, each one must be equally capable of being used to progress the operation.

A referral is not returned for a one-level or subtree search in which the search scope spans multiple naming contexts, and several different servers would need to be contacted to complete the operation. Instead, one or more search continuation references are returned. Search continuation references are intermixed with returned search entries. Each one contains a URL to another server (or set of servers) to contact, and represents an unexplored subtree of the namespace which potentially satisfies the search criteria. When multiple URLs are present in a given search continuation reference, each one must be equally capable of being used to progress the operation. By using a separate response for each unexplored subtree, LDAP Version 3 overcomes the limitation of LDAP Version 2, allowing alternate, equivalent servers to be included in each response.

As mentioned earlier, the other limitation in LDAP Version 2 referral processing is related to the inability of a client to specify whether a request was targeted for a normal object or a referral object. For LDAP Version 3, this difficulty is overcome with a protocol extension in the form of the manageDsaIT control. (Appendix G, "Supported server controls" on page 431 describes manageDsaIT in detail.) For typical client requests where the control is absent, whenever the server encounters an applicable referral object while processing the request, either a referral or search continuation reference is presented. When the client request includes this control, the server does not present any referrals or search continuation references, but instead treats the referral objects as normal objects. In this case, even superior references through the use of default referrals are suppressed. The shipped command line utilities support the -M option to indicate that the requestor is managing the namespace, and therefore wishes to examine and manipulate referral objects as if they were normal objects.

Example: associating servers through referrals and replication

Following are the steps involved in distributing the namespace using referrals.

1. Plan your namespace hierarchy.

```
country - US
company - IBM, Lotus
organizationalUnit - IBM Austin, IBM Endicott, IBM HQ
```

2. Set up multiple servers, each containing portions of the namespace.

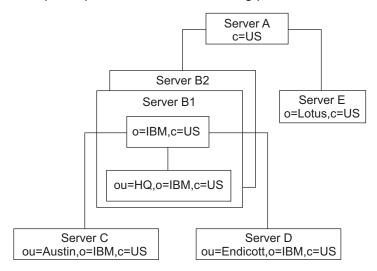


Figure 57. Setting up the servers

Following is a description of each server:

Server A

Perhaps just a server used to locate other servers in the US. With no other knowledge, clients can come here first to locate information for anyone in the US.

Server B1

A hub for all data pertaining to IBM in the US. Holds all HQ information directly. Holds all knowledge (referrals) of where other IBM data resides.

Server B2

A replica of Server B1.

Server C

Holds all IBM Austin information.

Server D

Holds all IBM Endicott information.

Server E

Holds all Lotus[™] information.

3. Set up referral objects to point to the descendents in other servers.

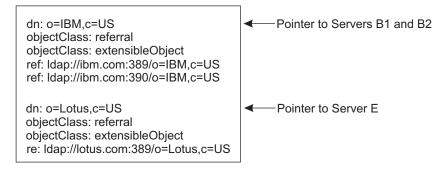


Figure 58. Server A database (LDIF input)

4. Servers can also define one or more default referrals which point to "more knowledgeable" servers for anything that is not underneath them in the namespace.

The default referrals go in the configuration file, not the backend.

Note: The default referral LDAP URLs do not include the DN portion.

General Section referral ldap://ibm.com:389 referral ldap://ibm.com:390 listen Idap://:789 # tdbm database definitions dabase tdbm "ou=Endicott,o=IBM,c=US" suffix

Figure 59. Server D configuration file

5. Putting it all together.

Figure 60 on page 283, Figure 61 on page 284, and Figure 62 on page 285 show these same six servers, showing the referral objects in the database as well as the default referrals which are used for superior references. Also included in Servers B1 and B2 are sample definitions for replication, setting up Server B as a replica of Server A. This ensures that these two servers will remain identical.

Server A: Services "c=US" **TDBM Database** dn: o=IBM,c=US objectClass: referral objectClass: extensibleObject ref: Idap://ibm.com:389/o=IBM,c=US ref: Idap://ibm.com:390/o=IBM,c=US ref: ldap://ibm.com:390/o=IBM,c=US dn: o=Lotus,c=US objectClass: referral objectClass: extensibleObject ref: Idap://lotus.com:389/o=Lotus,c=US

Server E: Services "o=Lotus,c=US" Configuration File referral Idap://US.white.pages.com:1234 Database dn: cn=Mikey,o=Lotus,c=US objectClass: person

Figure 60. Referral example summary (servers A and E)

Server B1: Services "o=IBM,c=US" Configuration File referral Idap://US.white.pages.com:1234 Listen Idap://:389 **TDBM Database** dn: cn=ReplicaB2,cn=localHost objectClass: replicaObject replicaHost: ibm.com replicaPort: 390 replicaBindDN: cn=Master replicaCredentials: secret dn: ou=Austin,o=IBM,c=US objectClass: referral objectClass: extensibleObject ref: Idap://austin.com:389/ou=Austin,o=IBM,c=US dn: ou=Endictt,o=IBM,c=US objectClass: referral objectClass: extensibleObject ref: Idap://endicott.com:789/ou=Endicott,o=IBM,c=US

Server B2: Services "o=IBM,c=US"

Configuration File

referral Idap://US.white.pages.com:1234 masterServer: Idap://ibm.com:389 masterServerDN: cn=Master secret masterServerPW:

Listen Idap://:390

TDBM Database

dn: ou=Austin,o=IBM,c=US objectClass: referral

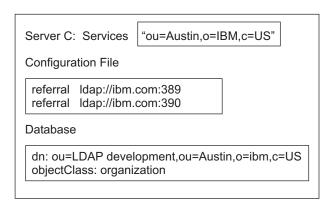
ref: Idap://austin.com:389/ou=Austin,o=IBM,c=US

dn: ou=Endictt,o=IBM,c=US

objectClass: referral

ref: Idap://endicott.com:789/ou=Endicott,o=IBM,c=US

Figure 61. Referral example summary (servers B1 and B2)



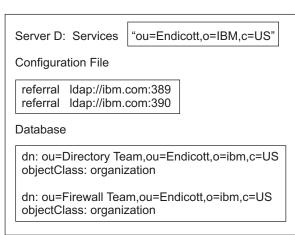


Figure 62. Referral example summary (servers C and D)

Chapter 26. Organizing the directory namespace

Directory services are meant to help organize the computing environment of the enterprise. To do this, directory services are meant to be used to help find all the resources at one's disposal. Information that is typically found in a directory consists of configuration information for services offered in the enterprise, locating information for people, places, and things in the enterprise, as well as descriptive information about services and resources available in the enterprise. The directory service should be thought of as the spot that can be queried to find whatever is desired in the enterprise.

When designing the format and organization of the directory service for an enterprise, the intended usage scenarios should be considered. These usage characteristics can have an impact on how the directory namespace should be organized so as to offer reasonable performance.

There are two general areas of directory namespace design to be considered. First, the types of information and the layout of where that information will be placed in the directory namespace must be determined. Additional information types can be added at a later date, but there should be some overall design of where in the directory namespace these types of information should be placed. Second, based on the usage characteristics of the users in the enterprise, the number of distinct directory servers and the namespace subtree or subtrees that they support must be considered.

As an example, consider an enterprise that consisted of two physical locations, one in Los Angeles, CA and one in New York City, NY. People in New York City access information about people, places, and things in Los Angeles often, while the people in Los Angeles rarely access information items in New York City. To offer good performance for both locations, a separate directory server could be installed and run in each location. These LDAP servers would manage information about the people, places, and things that reside in their respective locations. In addition, because the New York City personnel access information about things in the Los Angeles location, the information from the Los Angeles LDAP server could be replicated to an additional LDAP server at the New York City LDAP server. This would allow the New York City personnel to access information about the Los Angeles location by contacting a local server. In Los Angeles, however, directory requests about items in the New York City portion of the enterprise namespace are redirected (that is, referred) to the New York City LDAP server for the information. This would save managing a replicated set of information at the expense of slightly longer access times on the less-requested information.

The next two sections discuss information layout in the namespace and partitioning an enterprise namespace across multiple LDAP servers. These sections are followed by a small example.

Information layout

A directory is meant to provide information about people, places, and things in the enterprise. The most direct use of a directory is to hold information on how to contact other people in the enterprise. This has commonly been known as the *internal phone book*. With the widespread enhancements in technology, people are now more accessible than ever. We have pagers, answering machines, cellular phones, and e-mail. In trying to communicate with someone we might need to know about all of this information. Modeling a person object class based on the attributes about a person that are important to others in the enterprise is an easy way to support an online *internal phone book* using an LDAP directory service. In addition to people, different organizations within an enterprise can also be modeled by creating new object classes and attribute types. This would allow storage in the LDAP directory of locating information for useful services in the organization like benefits, travel reservations, and human resources.

Another application of directory services is the ability to model or store information about places. A place could be a conference room, which might have attributes of **numberOfSeats**, **projectorType**, **phoneNumber**, **calendarLocation**, **dataPortType**, **officeNumber**, and **buildingNumber**. Using this method, different conference rooms within a company could be located and compared. Another example of

a place would be the whole site in which employees work. An object class for a site LDAP directory entry might be made up of streetAddress, generalManagerDN, siteMap, and cafeteriaLocation.

Things abound within the enterprise. Under this category falls computers, copiers, FAX machines, printers, and computer software, as well as configuration information for servers that use an LDAP directory service. Each of these can be modeled with attribute types used inside object classes specific to the device or program.

In laying out where entries should appear in the directory hierarchy, by far the most common method of naming things is to start with the country in which the company is organized, followed by the name of the company, treated as an organization attribute type. Thus, the top level suffix for LDAP directory service names for entries within the company sometimes follows the form: o=CompanyName, c=US (for US-based companies). Alternately, the top level suffix may follow the domain form, for example: dc=CompanyName, dc=com. Below this suffix it is common for organizational unit object classes to be used to represent departments or sites within an organization. Below these organizational entries the actual entry representing a person, place, or thing would be defined. When organizing the information layout for the namespace, the intended usage should be considered to ensure the best performance.

Example of building an enterprise directory namespace

Let us look at an example configuration that exhibits the features available with the z/OS LDAP server. To set the stage, we will consider a moderately sized company that has personnel working in three locations across the United States. Big Company, Inc. has corporate headquarters in Chicago, IL, and two satellite facilities, one in Los Angeles, CA and the other in New York City, NY. The information technology staff would like to make available information about all of the company's computing and office services using an LDAP directory. In order to facilitate local modifications as necessary of the information in the directory, as well as provide improved response time for accessing local information, each site will have an LDAP server running. The server running at each site will be responsible for managing the directory information that pertains to that site.

The first thing to do is determine the name of the root of the directory namespace for Big Company, Inc. Typically, the name for the company will consist of the country of origin along with the company's given name. In LDAP directory terminology, the company is an organization. In this example, we chose: o=Big Company, c=US

as the company's name is Big Company and is located in the United States. Choosing a name of this format helps ensure that when a global namespace coordinator is established, the company's chosen root will fit nicely into the overall directory namespace.

Next to choose are the names of the three locations under which the directory information is stored. At this point, the namespace could be organized in a number of ways. One way would be to organize by functional unit (regardless of location). This model is useful if individuals (or computers, or other equipment or services) typically remain with the functional unit as opposed to being tied to the individual or physical location. Another way would be to organize based on the physical locations of the parts of the organization. This is useful if the people, places, and things to be stored in the directory typically do not move between locations. This latter approach will be used in the example. So, with three locations, three names are defined below the overall company distinguished name:

```
ou=Los Angeles, o=Big Company, c=US
ou=Chicago, o=Big Company, c=US
ou=New York City, o=Big Company, c=US
```

Since separate LDAP servers will be established at each location, these names represent the root of the subtree stored and managed by the directory server at each location.

For administration, each site will have a different directory administrator. To define this administrator, an administrator distinguished name and password need to be defined for each location. To start, the following names will be used:

```
AdminDN "cn=Administrator, ou=Los Angeles, o=Big Company, c=US"
AdminDN "cn=Administrator, ou=Chicago, o=Big Company, c=US"
AdminDN "cn=Administrator, ou=New York City, o=Big Company, c=US"
```

Since the Chicago location is also the corporate headquarters, the LDAP directory at this location will be used to store information about the entire company as well as information about the Chicago site.

We now have enough information to set up the base configuration files for each of the three LDAP servers that will be used to supply this information. Following are the files needed to set up the LDAP servers on each site. Note that what is shown is the minimal setup required. Other options could be specified in addition to these. See "Creating the slapd.conf file" on page 51 for configuration file options.

```
# Configuration file for the Chicago LDAP server
adminDN "cn=Administrator, ou=Chicago, o=Big Company, c=US"
database tdbm GLDBTDBM
suffix "o=Big Company, c=US"
dbuserid user1
databasename dbldap1
servername loc1
# end of configuration file
```

Figure 63. Chicago base configuration

```
# Configuration file for the Los Angeles LDAP server
referral ldap://ldap.chicago.bigcompany.com
adminDN "cn=Administrator, ou=Los Angeles, o=Big Company, c=US"
database tdbm GLDBTDBM
suffix "ou=Los Angeles, o=Big Company, c=US"
dbuserid user2
databasename dbldap2
servername loc1
# end of configuration file
```

Figure 64. Los Angeles base configuration

```
# Configuration file for the New York City LDAP server
referral ldap://ldap.chicago.bigcompany.com
adminDN "cn=Administrator, ou=New York City, o=Big Company, c=US"
database tdbm GLDBTDBM
suffix "ou=New York City, o=Big Company, c=US"
dbuserid user3
databasename dbldap3
servername loc1
# end of configuration file
```

Figure 65. New York City base configuration

The referral line indicates the default place to refer connecting clients when the LDAP server does not contain the information requested by the client. It is called the default referral. It is in the form of an LDAP URL. After the scheme name (1dap), the LDAP URL contains a TCP/IP DNS host name for another LDAP server. In this example, it is assumed that the TCP/IP host on which the Chicago LDAP server is running is Idap.chicago.bigcompany.com. The Chicago LDAP server does not have a default referral defined. This keeps directory searches from inadvertently going over the Internet from within the company.

The adminDN line indicates the distinguished name that should be used to connect to the LDAP server in order to have complete control over the data content held by the LDAP server.

The database line indicates that all following lines pertain to the TDBM storage method. The suffix line indicates what part of the namespace is contained in this server.

The next lines are used to connect to DB2 as well as identify the DB2 tables (named based on userid) that were defined for the LDAP server (see "Creating the DB2 database and table spaces for TDBM" on page 40).

After these files have been created and the SPUFI scripts have run successfully, one or more of the LDAP servers can be started. However, there will be no initial data in the DB2 tables. The next section introduces a load utility that can be used to load entries into the LDAP server.

Priming the directory servers with information

There are several methods of adding entries to an LDAP directory server: using the TDBM load facility (Idif2tdbm), using the Idapadd and Idapmodify tools, or using the LDAP C language API and the LDAP protocol. If you are not using **Idif2tdbm**, it is recommended that at least the top levels of directory information be loaded first into the database. This provides a base from which to add more entries into the directory namespace. If you are using **Idif2tdbm**, add the top levels and the additional entries at the same time to achieve the best load performance.

The Idif2tdbm utility program for TDBM can do the check and prepare phases while the LDAP server is running, but the actual loading of data cannot be done while the LDAP server is running. The load utility program makes direct updates to the directory database tables without using the LDAP protocol. This method results in faster loads of large amounts of directory information. The load utility takes as input a sequential file that contains the data describing directory entries to be added into the directory namespace. The format of the information in the sequential (HFS) file is LDAP Interchange Format (LDIF).

Using LDIF format to represent LDAP entries

The LDAP Data Interchange Format (LDIF) is used to represent LDAP entries in a simple text format. An LDIF file contains groups of attribute information which will be treated as an entry to be added to the directory. The general format of an LDIF entry is:

```
dn: distinguished name
attrtype1: attrvalue1
attrtype2: attrvalue2
```

Each line in the LDIF file must begin in column 1. However, to continue a line, start the next line with a single space or tab character. For example:

```
dn: ou=departments, ou=New York City, o=Big Co
mpany, c=US
```

Multiple attribute values are specified on separate lines. For example:

```
objectclass: organizationalunit
ou: departments
```

Note about editing LDIF files

Be aware that some editors, including oedit, place blank spaces at ends of all empty lines within a file. A blank space at the beginning of a line signifies continuation of the entry. The blank lines used to separate entries may be treated as continuations of an attribute value instead of separators if an editor has modified the LDIF file. Also, be aware that some editors, including oedit, delete blanks at the end of a line that is not empty. This can change the value of an attribute, especially if that value is continued on the next line.

If an attrvalue contains a nonprinting character, or begins with a space or a colon (:), the attrtype is followed by a double colon (::) and the value is encoded in base64 notation. For example, the value:

```
" begins with a space"
```

would be encoded like this:

```
cn:: IGJ1Z21ucyB3aXRoIGEgc3BhY2U=
```

Multiple entries within the same LDIF file are separated by blank lines. Here is an example of an LDIF file containing three entries.

```
dn: ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: New York City
dn: ou=fax machines, ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: fax machines
dn: ou=computers, ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: computers
```

Note: Trailing spaces are not trimmed from values in an LDIF file. Also, multiple internal spaces are not compressed. If you do not want them in your data, do not put them there.

Multiple attribute values for the same attribute type are specified on multiple lines within the specification of a directory entry. For example:

```
dn: cn=John Doe, ou=New York City, o=Big Company, c=US
objectclass: person
cn: John Doe
phonenumber: 555-1111
phonenumber: 555-2222
sn: Doe
```

Generating the file

A file is typically generated using an existing source of information and some tools to format the data into the LDIF format. Note that the order of entries in the LDIF file is important. In order for an entry specified in the LDIF file to be successfully added to the directory, its parent entry must first exist in the directory namespace. For this reason, the top level entries in the directory namespace subtree that the particular LDAP server will support must be first in the LDIF file.

For our example, we will define just the a minimal set of entries to get the directory server useful at each location. This will include two referral entries for the Chicago location. The meaning of these entries will be discussed in more detail in the following sections.

Here is the base set of LDIF files to set up the directory namespace at each location. For the Los Angeles location:

```
dn: ou=Los Angeles, o=Big Company, c=US
objectclass: organizationalunit
ou: Los Angeles
dn: cn=Administrator, ou=LosAngeles, o=Big Company, c=US
objectclass: person
cn: Administrator
sn: Administrator
userpassword: xxxxx
dn: ou=fax machines, ou=Los Angeles, o=Big Company, c=US
objectclass: organizationalunit
ou: fax machines
dn: ou=computers, ou=Los Angeles, o=Big Company, c=US
objectclass: organizationalunit
ou: computers
dn: ou=departments, ou=Los Angeles, o=Big Company, c=US
objectclass: organizationalunit
ou: departments
For the New York City location:
dn: ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: New York City
dn: cn=Administrator, ou=New York City, o=Big Company, c=US
objectclass: person
cn: Administrator
sn: Administrator
userpassword: xxxxx
dn: ou=fax machines, ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: fax machines
dn: ou=computers, ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: computers
dn: ou=departments, ou=New York City, o=Big Company, c=US
objectclass: organizationalunit
ou: departments
For the Chicago location:
dn: o=Big Company, c=US
objectclass: organization
o: Big Company
dn: ou=Los Angeles, o=Big Company, c=US
objectclass: referral
objectclass: extensibleObject
ref: ldap://ldap.losangeles.bigcompany.com/ou=Los Angeles,o=Big Company,c=US
dn: ou=New York City, o=Big Company, c=US
objectclass: referral
objectclass: extensibleObject
ref: ldap://ldap.newyorkcity.bigcompany.com/ou=New York City,o=Big Company,c=US
dn: ou=Chicago, o=Big Company, c=US
objectclass: organizationalunit
ou: Chicago
dn: cn=Administrator, ou=Chicago, o=Big Company, c=US
objectclass: person
```

```
cn: Administrator
sn: Administrator
userpassword: xxxxx
dn: ou=fax machines, ou=Chicago, o=Big Company, c=US
objectclass: organizationalunit
ou: fax machines
dn: ou=computers, ou=Chicago, o=Big Company, c=US
objectclass: organizationalunit
ou: computers
dn: ou=departments, ou=Chicago, o=Big Company, c=US
objectclass: organizationalunit
ou: departments
```

These files will now be used with a load facility. For a TDBM backend database, use the following command from the z/OS shell:

```
ldif2tdbm -o output.datasetHLQ -i ldif.filename -f config.filename -cpl
```

You can also run Idf2tdbm from TSO or from the batch environment.

After running this command on each respective directory server system, the directory namespace will be formed and the servers can now be used to hold and supply information.

Two entries added to the Chicago location directory server database deserve some special attention. These are the referral objects that were in the LDIF file for the Chicago location. Notice that the referral objects have the identical distinguished name as the root of the LDAP directory namespace that is served by the Los Angeles and New York City servers. These entries, coupled with the default referral specification in the configuration file for the LDAP servers in Los Angeles and New York City, enable searches of the Big Company namespace to originate at any of the three directory servers and resolve to the correct server to obtain the information.

A referral redirects a client request to a different LDAP server that can presumably handle the request (or refer the client to another server that can handle the request). In our example, if a client connects to the New York City server requesting a name that is under the Los Angeles portion of the namespace, the New York City server will send back a referral to the client based on the default referral. This will point the client at the Chicago directory server. The Chicago server will resolve the request down to the referral object for distinguished name ou=Los Angeles, o=Big Company, c=US and refer the client to the Los Angeles server. Finally, the client will contact the Los Angeles server and obtain the information requested.

Setting up for replication

As people start using the directory service in their daily routines at Big Company, Inc., the information technology staff notices that the people in New York City are doing a lot of work with the people in Los Angeles. So much, in fact, that an analysis of the TCP/IP traffic between New York City and Los Angeles shows that much of the traffic is directory access requests, presumably to look up phone numbers or FAX numbers for people in Los Angeles. The information technology staff decides to improve directory lookup response time, as well as lessen the directory lookup traffic between New York City and Los Angeles, by creating a replica of the Los Angeles directory server's information in New York City. This will allow local access to this information by the New York City users and cut down on the amount of requests from New York that must travel to Los Angeles to be completed.

Defining another LDAP server

To set up a replica of the LDAP server information in Los Angeles, a second LDAP server must be defined and started in New York City. This server can reside on the same system as the first LDAP server, though if this is chosen, the TCP/IP port that this replica server listens on must be different from the other LDAP server running on the system. As an alternative, the replica server could run on a different system,

allowing it to listen on the default LDAP port. The configuration file for the replica server in New York City will be very similar to the configuration files for the New York City server and the Los Angeles server. This configuration file must contain some additional items that pertain to replication. Here is what the contents of the New York City Los Angeles replica server should contain:

```
# Configuration file for the New York City Los Angeles replica LDAP server
referral ldap://ldap.chicago.bigcompany.com
listen ldap://:2001
adminDN "cn=Administrator, ou=Los Angeles, o=Big Company, c=US"
database tdbm GLDBTDBM
suffix "ou=Los Angeles, o=Big Company, c=US"
dbuserid user2
databasename dbldap
servername loc1
masterServer ldap://ldap.losangeles.bigcompany.com
masterServerDN "cn=Replicator, ou=Los Angeles, o=Big Company, c=US"
# end of configuration file
```

The additional lines at the end of the configuration file specify the only "user" that can update entries in the replica LDAP server. The values here must match the values entered at the "source" location when the replica is defined.

Preparing the replica

The next step is to get the LDAP replica primed with the existing information in the Los Angeles server and set up the Los Angeles server to replicate to the New York City replica. The set of steps to perform (described in "Populating a replica" on page 272) ensures that the replicas are in sync and that no update is lost during this synchronization. Once the replica is defined at the source location, updates to the directory information will be logged to be sent to the replica server when possible.

To initially synchronize the data between the LDAP master server and the LDAP replica server, perform the steps in "Populating a replica" on page 272.

While there are a number of manual steps to perform, there is a small consolation that the steps at different locations are not interleaved. All work can be done at the source location and then all work can be performed at the target (replica) location.

Resynching the replica and master servers

If it is noticed that a replica's contents are out of sync with the information at the master server, the information can be resynched by following the steps shown in "Recovering from out-of-sync conditions" on page 275.

Notifying users of the replica

At this point, the New York City users can be notified that a second LDAP server is now available for their use. The notification should contain either the LDAP URL of the new LDAP replica server or the host name and port number of the LDAP replica server, as well as the base of the LDAP subtree that is published by the replica. As updates are made to the Los Angeles LDAP server, these updates will be propagated to the replica server in New York City. See Chapter 24, "Replication" on page 269 for more details on replication.

What Big Company, Inc. now has in place is an Enterprise Directory service that can be used by whatever enterprise distributed processing tasks require lookup or configuration information. These enterprise distributed processing tasks and applications may require some changes to make use of the directory

service, but the result will be the ability to view, find, and modify the configuration of the enterprise by looking at and modifying the contents of the LDAP directory.

Chapter 27. Client considerations

When an LDAP application is communicating with an z/OS LDAP server, you should consider the following special topics:

- · Root DSE
- Monitor Support
 - UTF-8 data over the LDAP Version 2 protocol
 - Attribute types stored and retrieved in lowercase
 - · Reason codes
- CRAM-MD5 authentication support

Root DSE

Following is an example of the information that the z/OS LDAP server will report on a search of the root DSE. A subset of these values may appear in your rootDES based on the server configuration choices you have made.

```
namingcontexts=cn=localhost
namingcontexts=o=IBM,c=US
altserver=ldap://host2.ibm.com:999
altserver=ldap://host3.ibm.com:999
ref=ldap://hostk.ibm.com:391
ref=ldap://host1.ibm.com:333
supportedsas1mechanisms=CRAM-MD5
supportedsas1mechanisms=DIGEST-MD5
supportedsas1mechanisms=EXTERNAL
supportedsas1mechanisms=GSSAPI
ldapservicename=host1.ibm.com@IBM.COM
supported1dapversion=2
supported1dapversion=3
supportedcontrol=2.16.840.1.113730.3.4.2
supportedcontrol=1.3.18.0.2.10.2
supportedcontrol=1.3.18.0.2.10.6
supportedcontrol=1.3.18.0.2.10.10
supportedcontrol=1.3.18.0.2.10.11
supportedextension=1.3.6.1.4.1.1466.20037
supportedextension=1.3.18.0.2.12.8
supportedextension=1.3.18.0.2.12.7
ibmdirectoryversion=z/OS V1R4
ibm-sasldigestrealmname=host1.ibm.com
```

Monitor Support

Monitor support allows customers or applications to search the server with a base of *cn=monitor* and have the server return some statistics regarding the operation of the server. Following is an example:

```
ldapsearch -h dceset3 -p 389 -b cn=monitor -s base objectclass=*
cn=monitor
version=z/OS Version 1 Release 4 Security Server LDAP Server
totalconnections=1
currentconnections=1
livethreads=10
maxconnections=60
opsinitiated=2
opscompleted=1
entriessent=0
bytessent=0
currenttime=Thu Oct 18 15:40:06 2001 EDT
starttime=Thu Oct 18 15:09:42 2001 EDT
```

CRAM-MD5 Authentication Support

CRAM-MD5 authentication is supported on the IBM Directory Server and client utilities. However, the way that it has been implemented on IBM Directory Server is different than the z/OS LDAP server. Thus, this has resulted in differences between the IBM Directory Server and the z/OS LDAP server client utilities. In order to perform a CRAM-MD5 authentication bind with the IBM Directory Server client utilities to the z/OS LDAP server, you must specify the bind DN (authorization DN) with the -D option. The IBM Directory Server client utilities do not support specifying an authentication identity (username). See Chapter 21, "CRAM-MD5 and DIGEST-MD5 Authentication" on page 243 for more information on CRAM-MD5 and DIGEST-MD5 bind authentication on the z/OS LDAP server.

UTF-8 data over the LDAP Version 2 protocol

The LDAP Version 3 Protocol allows UTF-8 attribute values outside of the IA5 character set to be stored in the directory. This information must be able to be returned in some format to LDAP Version 2 clients. An additional server configuration file option, sendV3stringsoverV2as, which has the possible values ISO8859-1 or UTF-8, can be used to indicate which format to use when sending this information over the Version 2 protocol.

Note: Different clients treat non-IA5 data differently over the Version 2 protocol. Refer to the documentation for the client APIs you are using for more information.

Attribute types stored and returned in lowercase

The LDAP server stores and returns attribute types in lowercase (normalized). For example, the attribute type "productName" is returned as "productname".

Abandon behavior

In releases of the LDAP server prior to z/OS V1R4, an abandon operation did not affect the processing of the targeted operation because all operations were handled by the server in a sequential order. By the time the abandon was received, the operation that was the target of the abandon had already completed.

In z/OS V1R4, the LDAP server reads additional operations as they arrive as long as the connection is not a secure connection and the previous operation is not bind, unbind, or extended operation. This allows the LDAP server to process abandon operations as they are received and affect previously submitted operations.

Changed reason codes

The return code provided by the LDAP server has changed from LDAP PROTOCOL ERROR to LDAP UNWILLING TO PERFORM for several errors related to update of the LDAP server's TDBM schema. This change was made for compatability purposes with LDAP servers on other IBM platforms. These errors include:

- The schema is missing a definition for an attribute that is defined as required by an objectclass definition.
- The schema is missing a definition for an attribute that is defined as optional by an objectclass definition.
- The schema has a duplicate definition of an attribute that is defined as required by an objectclass definition.
- The schema has a duplicate definition of an attribute that is defined as optional by an objectclass definition.
- The schema is missing a definition for an objectclass that is in the superior chain of another objectclass.
- · A cycle has been detected in the superior chain of an objectclass definition.

• The objectclass top is not defined in the superior list of a STRUCTURAL objectclass.

The return code for attempting to bind to a DN that does not exist in any of the backends condfigured on the LDAP server has been changed from LDAP_NO_SUCH_OBJECT to LDAP_INVALID_CREDENTIALS. The reason this return code has been changed is to provide better security in our LDAP server since it prevents discovery of a valid DN.

Reason codes

The LDAPResult construct is used by the LDAP protocol to return success or failure indications from servers to clients. This construct contains an error message field. Servers can optionally provide "human-readable" diagnostic information in this field. Depending on the location in the z/OS LDAP server where errors are detected, error messages generated may have the following format:

R<numeric digits> <diagnostic information> <traceback information>

where:

numeric digits

Represents a specific reason code.

diagnostic information

Provides details about the reason for the failure.

traceback information

Is of the form (file_identification | file_version | line_number) and will assist you in diagnosing application or configuration problems.

Note the following regarding this error information:

- It is intended to be "human-readable" to assist in identifying problems detected by the server.
- It is not translated (English text only).
- It is not intended to be used as an application programming interface (API).
- · Data returned may be changed by service or new releases of the product. (Again, it is not intended to be an API.)

Following is the current list of reason codes and associated diagnostic information returned by the z/OS LDAP server.

R000001	An attempt to allocate heap storage failed.	R000008	Error on BER encoding of attribute attr_type and its values.
R000002	An attempt to allocate heap storage failed. Length: length.	R000009	Encoding error on BER encoding value(s) of attribute attr_type.
R000004	An internal server error has been encountered.	R000010	The filter used to search for the schema is not valid. Filter should be objectclass=subschema.
R000005	Attempt to translate data from source_codepage to target_codepage failed with LDAP rc=0xrc.	R000011	Error on BER decoding of attribute(s): attribute-name
R000006	Error on BER encoding of BER_element.	R000100	The password has expired.
R000007	Error on BER encoding of attribute	R000101	The new password is not valid.
	attr type.		

R000102	The user ID has been revoked.	R000120	Cannot specify more than one value for attribute attr_type.
R000103	The user ID is not defined.	R000121	Cannot specify a value for attribute
R000104	The password is not correct or the userid is not completely defined (missing password or uid).		attr_type that is different than the value in the DN.
		R000122	The value for attribute attr_type must be the DN of a user.
R000105	A bind argument is not valid.		be the DN of a user.
R000106	An unknown error occurred.	R000123	The value for attribute attr_type must be the DN of a group.
R000107	SDBMBeCapability: Constructor missing a required initializer.	R000124	The value for attribute attr_type must be the DN of a user or a group.
R000108	Constructor call to entry_attr_merge_becap failed.	R000125	Attribute attr_type is not supported.
R000109	SDBMBeCapability::operator= not implemented.	R000126	Filter filter is not supported for this base.
R000110	SDBMBeCapability::operator== not implemented.	R000127	Filter filter contains a type without a value.
R000111	SDBMBeCapability::default constructor not implemented.	R000128	Filter filter is not supported.
	not implemented.	R000129	Value value is not supported for filter
R000112	SDBMBeCapability::getNextVal: Bad cursor argument.		filter.
		R000130	The syntax of filter filter is not valid.
R000113	SDBMBeCapability::getNextVal: NULL cursor argument.	R000131	DN is not a valid RACF DN.
R000114	The realm portion of the value of attribute attr_type is not the RACF default realm.	R000132	Value for attribute attr_type cannot be more than length characters.
		R000133	Value for attribute attr_type must be an
R000115	There is no RACF default realm.		integer less than size.
R000116	Cannot specify a value when deleting attribute attr_type.	R000134	The RACF command command created to satisfy this request is too long, probably due to specifying a long filter
R000117	Cannot delete attribute attr_type.		or attribute value or too many attribute values.
R000118	Cannot replace attribute attr_type.	R000135	Cannot perform this request on a reserved RACF DN, DN.
R000119	Cannot add or replace attribute attr_type.		

R000136	Operation not supported. The schema entry cannot be used in the compare	R001015	Cycle detected in SUP_hierarchy.
	operation.	R001016	required_information missing for schema value: value.
R000137	dn is not a valid RACF DN for bind.		
	Check that the syntax is correct and that it is a DN for a RACF user.	R001017	Syntax/matching rule inconsistency for attribute type: attr_type.
R001001	value is not a valid Generalized Time format value.	R001018	Value specified is obsolete. Value:value.
R001002	Initialization of required schema failed.	R001019	Internal error during schema
R001003	RDN specified for schema entry not		azao
	valid.	R001020	User schema did not contain required valuevalue.
R001004	Schema entry contained unknown		
	information.	R001021	User schema did not contain required value value for value.
R001005	Duplicate value encountered: value.		
R001006	Value not found in schema. Value:value.	R001022	User and required schema keyword values did not match for attr_type.
		R001023	Building minimum schema failed.
R001007	Program used improperly set pointer.		
R001008	Specified value did not match syntax.	R001024	Abstract object class: object_class can not be leaf.
R001009	Specified value did not match syntax. Value: value.	R001025	Multiple structural object classes found as leaf:object_class, object_class.
R001010	Specified value did not match syntax. Value: <i>value</i> .	R001026	Structural object class not found for entry.
R001011	Collective keyword not supported for attribute type: attr_type.	R001027	Entry attempted to change base structural object class from: object_class to object_class.
R001012	Attribute information not found		
	for:attr_type.	R001028	More than one value specified for single valued attribute type: attr_type.
R001013	Usage designated for superior attribute	D004000	Most attained to a net formal in
	type not consistent with attribute: attr_type.	R001029	Must attribute type not found in entry: attr_type.
R001014	Type designated for superior object class: object_class not consistent with object class: object_class.	R001030	Entry contained attribute type not allowed by schema: attr_type.

R001031	Missing left parenthesis in schema value: value.
R001032	Missing right parenthesis in schema value: value.
R001033	Schema processing encountered an internal error.
R001034	An unsupported option was found for keyword: keyword in value: value.
R001035	End of data found when expecting more values in value:value.
R001036	Length value was not specified in value: value.
R001037	Non-numeric length specified in value: value.
R001038	Format of numeric object identifier not valid in value: value.
R001039	Unsupported non-alphabetic character 'character' at offset offset in value: value.
R001040	Unsupported character 'character' at offset offset in value: value.
R001041	Values not found where expected in value: value.
R001042	Missing a leading delimiter (delimiter) in a portion of the value: value.
R001043	Missing a trailing delimiter (delimiter) in a portion of the value: value.
R001044	A required blank is missing at offset offset in value:value.
R001045	Keyword: keyword not supported for value: value.
R001046	Missing closing quote for value: value.

R001047	Missing leading quote for value: value.
R001048	Missing closing brace for value: value.
R001052	Non-numeric characters found in an Integer value: value.
R001053	Integer value of length length exceeds the maximum allowable size (size).
R001054	DirEntry parameter initialized on input.
R001055	Specified attribute: attr_value not valid for schema entry.
R001056	Specified object class: object_class not valid for schema entry.
R001057	Schema entry missing required object class.
R001058	Schema entry missing required value.
R001059	User schema did not contain required value value for value.
R001060	Superior object class: object_class found during sups processing is obsolete.
R001061	The character 'character' (0xhex_character) is not supported in substring search filters of attribute type:attr_type.
R001062	Substring search filters are not supported for attribute type: attr_type.
R001063	The keyword 'keyword' is specified multiple times in value:value
R001064	'date/time' specifier is not valid for value: value. Valid values arevalues.
R001065	Attribute information not found for: attr_type.

nema did not contain required alue for value.	R002009 R002010 R002011	Keyword: keyword not allowed multiple times in: value. Search filter processing failed calling regcomp(), rc=rc("error_string"), pattern="pattern". Search filter processing failed calling regexec(), rc=rc("error_string"), comparison value="value".
attribute type: attr_type failed, susage is not userApplications. ce attribute type not found arsing IBM Attribute	R002011	regcomp(), rc=rc("error_string"), pattern="pattern". Search filter processing failed calling regexec(), rc=rc("error_string"),
c usage is not userApplications. ce attribute type not found ursing IBM Attribute		<pre>regexec(), rc=rc("error_string"),</pre>
rsing IBM Attribute	 R002012	comparison value="value".
_type.	RUUZUIZ	LDIE line missing congretor (%)
		LDIF line missing separator (':') between attribute type and value: value.
entry did not contain all attribute types.	R002013	LDIF line contains an attribute type (attr_type), but no value.
d not contain an object class.	R002014	Incorrect base-64 encoded LDIF value
an one object class type I found in value: <i>value</i>		found. Attribute type=attr_type. Bad character=character' (Oxhex_character).
equal sign in DN ent:DN_component.	R002015	LDIF entry contains multiple DNs (first DN in entry= <i>DN</i>).
pe sequence found in the g DN component is not // component.	R002016	The following LDIF entry does not contain a DN: <newline>LDIF entry.</newline>
nal server error has been ered.	R002017	The DN component DN_component contains an unescaped special character (char).
closing quote, or incomplete sequence in the DN ent: DN_component.	R002018	An extraneous colon was found in aclEntry value: value.
cter formed by an escape	R003001	ACLManager::checkAccess() is not implemented.
ee in the following DN ent is not valid UTF-8: ponent.	R003002	ACLManager::applyAdd() is not implemented.
ON components are not ed.	R003003	ACLManager::applyAddForSrcEID() is not implemented.
of aclEntry attribute not as d:value.	R003004	ACLManager::applyAddForAclsrcEID() is not implemented.
ions missing for security level	R003005	ACLManager::applyAddForOwnsrcEID() is not implemented.
	ent is not valid UTF-8: conent. N components are not d. f aclEntry attribute not as evalue.	ent is not valid UTF-8: Ponent. R003002 R003003 R003003 R003004 R003004

R003006	ACLManager::applyAddForPropagate() is not implemented.	R003022	ACLManagerFactory::deleteACLObject() is not implemented.
R003007	ACLManager::applyAddModifyPreScan() is not implemented.	R003023	ACLManagerFactory::deleteACLSubject() is not implemented.
R003008	ACLManager::applyAddModifyPost Apply() is not implemented.	R003024	ACLManagerFactory::deleteACLRights() is not implemented.
R003009	ACLManager::applyAddModify() is not implemented.	R003025	ACLManagerFactory::deleteACLMod Progress() is not implemented.
R003010	ACLManager::applyAddDelete() is not implemented.	R003026	Failure in constructing the ACLManager instance, rc =rc.
R003011	ACLManagerFactory::newACLManager() is not implemented.	R003027	Unable to determine parent EID during set of ACLSRC EID.
R003012	ACLManagerFactory::newACLSubject() is not implemented.	R003028	Unable to find first propagating EID during set of ACLSRC EID.
R003013	ACLManagerFactory::newACLObject(1) is not implemented.	R003029	Either aclPropagate or aclEntry was missing from new entry.
R003014	ACLManagerFactory::newACLObject(2) is not implemented.	R003030	Unable to determine parent EID during set of OWNSRC EID.
R003015	ACLManagerFactory::newACLObject(3) is not implemented.	R003031	Unable to find first propagating EID during set of OWNSRC EID.
R003016	ACLManagerFactory::newACLObject(4) is not implemented.	R003032	Either ownerPropagate or entryOwner was missing from new entry.
R003017	ACLManagerFactory::newACLObject(5) is not implemented.	R003033	Incorrect value specified for aclPropagate, value =value.
R003018	ACLManagerFactory::newACLObject(6) is not implemented.	R003034	Multiple values specified for aclPropagate.
R003019	ACLManagerFactory::newACLRights() is not implemented.	R003035	Incorrect value specified for ownerPropagate, value =value.
R003020	ACLManagerFactory::newACLMod Progress() is not implemented.	R003036	Multiple values specified for ownerPropagate.
R003021	ACLManagerFactory::deleteACLManager() is not implemented.	R003037	Incorrect aclPropagate setting found.

Incorrect aclPropagate setting found.	R003054	Access allowed for add because subject is the owner of the entry.
Either aclPropagate or aclEntry was missing from new entry during propagation processing.	R003055	Unable to find parent ACL during add check.
Either ownerPropagate or entryOwner was missing from new entry during propagation processing.	R003056	Access allowed because subject has add permission in parent ACL.
ACLManagerImpl::applyModify PreScan() is not implemented.	R003057	Access denied because subject does not have add permission in parent ACL.
ACLManagerImpl::applyModify PostApply() is not implemented.	R003058	Unable to find first propagating ACL EID during add check.
ACLManagerImpl::applyModify() is not implemented.	R003059	Unable to find effective ACL entry in first propagating ACL during add check.
Incorrect value =value found for aclprop.	R003060	Access denied for add because subject does not have an entry in propagating ACL.
Incorrect value =value found for ownprop.	R003061	Access allowed because subject has
Incorrect value specified for aclPropagate in modifications list,		write permission to all attributes in new entry.
Incorrect value specified for ownerPropagate in modifications list,	R003062	Access denied for add because subject does not have write permission to all attributes in new entry.
	R003063	Incorrect return code =rc encountered at end of access check.
Unable to determine if subject is the Administrator.	R003064	Unknown ACLRIGHTS value encountered.
Access allowed because subject is the Administrator.	R003065	Unable to determine if subject is an Owner during modify check.
Unable to build parent entry during check for add rights.	R003066	Access allowed for modify because subject is the owner of the entry.
Unable to find parent EID and parent DN is not a suffix.	R003067	Unable to find effective ACL entry in during modify check.
Unable to determine if subject is an Owner during add check.	R003068	Access denied for modify because subject does not have an entry in ACL.
	Either aclPropagate or aclEntry was missing from new entry during propagation processing. Either ownerPropagate or entryOwner was missing from new entry during propagation processing. ACLManagerImpl::applyModify PreScan() is not implemented. ACLManagerImpl::applyModify PostApply() is not implemented. ACLManagerImpl::applyModify() is not implemented. Incorrect value =value found for aclprop. Incorrect value specified for aclPropagate in modifications list, value =value. Incorrect value specified for ownerPropagate in modifications list, value =value. Unable to rebuild caches. Unable to determine if subject is the Administrator. Access allowed because subject is the Administrator. Unable to build parent entry during check for add rights. Unable to find parent EID and parent DN is not a suffix.	Either aclPropagate or aclEntry was missing from new entry during propagation processing. Either ownerPropagate or entryOwner was missing from new entry during propagation processing. ACLManagerImpl::applyModify PreScan() is not implemented. ACLManagerImpl::applyModify PostApply() is not implemented. ACLManagerImpl::applyModify() is not implemented. Incorrect value = value found for aclprop. R003059 Incorrect value = value found for ownprop. R003060 Incorrect value specified for aclPropagate in modifications list, value = value. R003062 Incorrect value specified for ownerPropagate in modifications list, value = value. Unable to rebuild caches. Unable to determine if subject is the Administrator. Access allowed because subject is the Administrator. Unable to build parent entry during check for add rights. Unable to find parent EID and parent DN is not a suffix.

R003069	Access allowed for modify because subject has write permission to all modified attributes.	R003083	Unable to determine if subject is an Owner during search check.
R003070	Access denied for modify because subject does not have write permission	R003084	Access allowed for search because subject is the owner of the entry.
R003071	to all modified attributes.	R003085	Unable to find effective ACL entry in during search check.
	Unable to determine if subject is an Owner during delete check.	R003086	Access denied for search because subject does not have an entry in ACL.
R003072	Access allowed for delete because subject is the owner of the entry.	R003087	Access allowed for search because
R003073	Unable to find effective ACL entry in during delete check.		subject has read permission on all requested attributes.
R003074	Access denied for delete because subject does not have an entry in ACL.	R003088	Access allowed for search because subject has read permission some requested attributes.
R003075	Access allowed for delete because subject has delete permission for the entry.	R003089	Access denied for search because subject does not have read permission on all requested attributes.
R003076	Access denied for delete because subject does not have delete permission on the entry.	R003090	Unable to determine if subject is an Owner during compare check.
R003077	Unable to determine if subject is an Owner during modify name check.	R003091	Access allowed for compare because subject is the owner of the entry.
R003078	Access allowed for modify name because subject is the owner of the	R003092	Unable to find effective ACL entry in during compare check.
	entry.	R003093	Access denied for compare because subject does not have an entry in ACL.
R003079	Unable to find effective ACL entry in during modify name check.	R003094	Access allowed for compare because subject has compare permission on
R003080	Access denied for modify name because subject does not have an		attribute.
 R003081	entry in ACL. Access allowed for modify name	R003095	Access denied for compare because subject does not have compare permission on attribute.
	because subject has write permission on all attributes in the name.	R003096	Access check indicated successful but
R003082	Access denied for modify name		allowed flag was not set.
	because subject does not have write permission on all attributes in the name.	R003097	Update of SRC column failed for case oldVal=NONE, newVal=TRUE.

R003098	Update of SRC column failed for case oldVal=NONE, newVal=FALSE.	R003113	ACLManager Factory::newACLobject(7) is not implemented.
R003099	Update of SRC column failed for case oldVal=TRUE, newVal=NONE.	R003114	Access denied because newSuperior could not be determined.
R003100	Update of SRC column failed for case oldVal=TRUE, newVal=FALSE.	R003115	Access denied for delete because subject does not have an entry in ACL.
R003101	Update of SRC column failed for case oldVal=FALSE, newVal=NONE.	R003116	Unable to determine if subject is an Owner during modify name check.
R003102	Update of SRC column failed for case oldVal=FALSE, newVal=TRUE.	R003117	Unable to find effective ACL entry during modify name check.
R003103	Update results in aclPropagate value without an associated aclEntry value.	R003118	Access denied for modify name (with newSuperior) because subject does not have an entry in ACL for target
R003104	Update results in no aclPropagate value with an associated aclEntry		object.
	value.	R003119	Access allowed for modify name because subject has write permission
R003105	Update results in ownerPropagate value without an associated entryOwner value.		on all attributes in the target object's name.
	entryowner value.	R003120	Access denied for modify name because subject does not have write permission on all atrributes in the name.
R003106	Update results in no ownerPropagate value with an associated entryOwner value.		
R003107	Unable to add ACL or owner information into entry.	R003121	Unable to build parent entry during check for modify DN rights.
R003108	ACLManager::isACLModified() is not implemented.	R003122	newSuperior designated as root and renamed DN is not a suffix (Modify DN).
R003109	Unable to determine first propagating Owner for new entry.	R003123	Access allowed for modify name because subject has necessary permissions on both target and
R003110	Unable to determine if subject would be Owner for new entry.		newSuperior objects.
R003111	ACLManager::getFirstPropagation AclEID() is not implemented.	R003124 	Access denied for modify name because subject does not have add permission on newSuperior object.
R003112	ACLManager::getFirstPropogating OwnEID() is not implemented.	R003125	Access denied for modify name (with newSuperior) because subject does not have delete permission on target object.

R003126	Access denied for modify name (with newSuperior) because subject does not have add permission on	R004016	Error decoding long attribute values for attribute type attr_type in entry DN.
	newSuperior object.	R004017	No attribute values found for entry DN.
R003127	Access denied for modify name (with newSuperior) because subject does not have an entry in ACL for newSuperior object.	R004018	Internal error encountered with the TDBMCacheManager.
		R004019	Entry data is missing required RDN
R004001	Unknown error occurred!		components.
R004002	Container operation failed!	R004020	RDN contains duplicate values for attribute attr_type, value = value.
R004003	The base dn <i>DN</i> specified on the command is not valid.	R004022	No parent entry for <i>DN</i> .
R004004	Encrypt of directory entry password failed with $rc = rc$.	R004023	Unable to create an ACL request object.
R004005	DN DN exceeds maximum length of max_length.	R004024	The filter used to search for the schema is not valid. Filter should be "objectclass=subschema".
R004006	DN DN already exists.	R004025	Error in search checking to see if DN is a V2 referral.
R004007	The parent DN DN specified is not		
	valid.	R004026	Entry not found in the database.
R004008	Object DN does not exist.	R004027	LDAP Search is unwilling to perform the attempted search. The generated
R004009	Data Base utilities failed with $rc = rc$.		SQL statement is too large.
R004010	Get ancestors utilities failed with rc = rc.	R004028	The size limit for your search has been reached.
R004011	SQL utilities failed with rc = rc.	R004029	Unable to send search results.
R004012	Error decoding DN from database.	R004030	Search failed because it was unable to check ACLs for returned entries.
R004013	Error decoding attribute type from		
	database for entry DN.	R004031	Time limit exceeded for the present search.
R004014	Error skipping the decoding of attribute type attr_type in entry DN.	R004032	Attempt to scan referral cache failed.
R004015	Error decoding attribute values for attribute type attr_type in entry DN.	R004033	Unrecognized filter type.

R004034	Unable to create entry structure for the entry object.	R004049	Attempt to delete entry <i>DN</i> from replica servers failed. Deletion was unsuccessful.
R004035	Attribute attr_type is not allowed to be		
	modified by the user.	R004050	Unable to create a request structure.
R004036	<pre>attr_type attribute already contains value = value.</pre>	R004051	Entry DN does not contain attribute attr_type.
R004037	If the attr_type attribute is specified its value must be equal to the DN.	R004052	Entry size column is missing from select statement columns.
R004038	Operation not allowed. The configuration file specifies a read-only mode for this server.	R004053	EID column is missing from select statement columns.
R004039	Operation not allowed. The schema cannot be deleted from the server.	R004054	One or more characters found in string string are not valid UTF-8.
R004040	An error occurred checking the referral cache.	R004055	Attribute attr_type is not allowed to be modified by the user.
R004041	Entry <i>DN</i> is not a leaf. Deletion is not allowed.	R004056	Modrdn of the schema entry is not allowed.
	Hardele to determine if out to DAVie	R004057	DBXGetData Failed.
R004042	Unable to determine if entry <i>DN</i> is a leaf. Deletion was unsuccessful.	R004058	Unable to get encryption method.
R004043	Unable to create SQL to delete from the descendants table. Deletion of entry <i>DN</i> is unsuccessful.	R004059	Credential cannot be encrypted.
	onaly 27010 and accession	R004060	Entry does not contain a password.
R004044	Unable to delete entry <i>DN</i> from the descendants table. Deletion was unsuccessful.	R004061	DBXExecDirect failed.
R004045	Unable to create SQL to delete from	R004062	Credentials are not valid.
1100-10-10	the search table. Deletion of entry <i>DN</i> is unsuccessful.	R004063	Unable to gather group information.
R004046	Unable to delete entry DN from the	R004064	DBXBindParameter failed.
	search table. Deletion was unsuccessful.	R004065	DBXFetch failed.
R004047	Unable to create SQL to delete from the DIR_ENTRY table. Deletion of entry	R004066	Unable to add/delete/replace values.
	DN is unsuccessful.	R004067	Authentication method is not SIMPLE.
R004048	Unable to delete entry <i>DN</i> from the DIR_ENTRY table. Deletion was unsuccessful.		

R004068	Modifications were not specified.	R004088	Entry DN does not contain attribute attr_type.
R004069	Server unable to obtain a request structure.	R004089	Attribute values not specified for replace.
R004070	Referral cache scan failed.	 R004090	Non UTF-8 V3 data detected.
R004071	DN does not exist.	R004091	IA5 fence is on and non-IA5 V2 data is
R004072	Unable to update the entry object.	1004031	detected.
R004073	Non-Leaf Referral test failed.	R004092	Error Converting V2 string from ISO8859-1 to UTF-8.
R004074	Unable to prepare for entry table updates.	R004093	Error converting binary data to a textual value.
R004075	Unable to remove old RDN values.	R004094	Attribute attr_type cannot be a component of the RDN.
R004076	DNS style names are not allowed.		·
R004077	DN already exists.	R004095	Filter contained a NOT with an unrecognized attribute type, which is unsupported.
R004078	Unable to prepare for search table updates.	R004096	Entry DN does not contain attribute attr_type, value=value.
R004079	Unable to perform database operations.	R004097	Error decoding the entry from the database.
R004080	User canceled request.	 R004098	Filtering on non-textual attributes
R004081	Unable to update the entry table.	1100-1000	attr_type is not allowed.
R004082	Replication updates failed.	R004099	Parent of new entry <i>DN</i> is referral entry <i>DN</i> . Operation not allowed.
R004083	New superior is not allowed.	R004100	TDBMBeCapability: Constructor missing a required initializer.
R004084	Entry is not a leaf.		
R004085	Out of memory.	R004101	Constructor call to entry_attr_merge_becap failed.
R004086	Entry DN already contains attribute attr_type, value=value.	R004102	TDBMBeCapability::operator= not implemented.
R004087	Encryption of password failed.	R004103	TDBMBeCapability::operator== not implemented.

R004104	TDBMBeCapability::default constructor not implemented.	R004119	Modify-add of the new password must occur after the modify-delete of the old password for native authentication	
R004105	TDBMBeCapability::getNextVal: Bad cursor argument.		password updates on entry <i>DN</i> .	
		R004120	Entry DN participates in native	
R004106	TDBMBeCapability::getNextVal: NULL cursor argument.		authentication so adding the userpassword attribute is not allowed.	
R004107	Thepasswd function failed; not loaded from a program controlled library.	R004121	Entry <i>DN</i> participates in native authentication but is not configured correctly. If useNativeAuth = ALL then the entry must contain the attribute ibm-nativeld or uid.	
R004108	TDBM backendpasswd API resulted in an internal error.	R004122	Unable to get cursor name for DELETE	
R004109	The password has expired.		statement.	
R004110	The user id has been revoked.	R004123	Allocation of statement handle failed.	
R004111	The password is not correct.	R004124	DELETE at cursor statement execution failed.	
	<u> </u>			
R004112 R004113	A bind argument is not valid. Native authentication cannot be	R004125	Operation not supported. The schema entry cannot be used in the compare operation.	
K004113	performed when multiple UID values		TI D	
	exist for entry DN.	R004126	The Prepare of a SQL statement failed.	
R004114	Modify-delete of the old password for entry <i>DN</i> must occur before the	R004127	The Execute of a SQL statement failed.	
	modify-add of the new password for native authentication password updates.	R004128 	Native authentication password change failed: \n. The new password is not valid, or does not meet requirements.	
R004115	More than one password cannot be specified for a native authentication password update on entry <i>DN</i> .	R004129	Unable to locate newSuperior <i>dn</i> in any backend.	
R004116	Password change not allowed for entry <i>DN</i> . The nativeUpdateAllowed configuration option has not been enabled.	R004130	Time limit exceeded for Modify DN operation - routine-name.	
		R004131	Determination of descendants to rename failed with rc = return-code.	
R004117	Native authentication password replace is not allowed for entry <i>DN</i> .	R004132 	Either the newSuperior DN must exist in this backend, or (if it does not) the new (renamed) DN MUST be a suffix in this backend.	
R004118	Entry <i>DN</i> native user ID (ibm-nativeld,uid) is not defined to the Security Server.			

R00)4133	The newSuperior DN is located in the subtree to be moved - not permitted.	R004148	An attribute type which should be present in candidate for DN realignment was not found.
R00)4134	Failure occurred while inserting new ancestor rows in hierarchy.	R004149	Unable to complete DN realignment - could not update entry object.
R00)4135	Failure occurred while deleting old ancestors rows from hierarchy.	R004150	Unable to complete DN realignment - could not update DIR_SEARCH with
R00	04136	Failure occurred while updating parent EID to newSuperior EID.		changes.
R00)4137	Failure occurred while updating subtree node levels.	R004151 	Failure occurred while attempting to lock candidate rows in DIR_ENTRY for modify DN operation.
R00)4138	Unable to locate effective ACLEntry for renamed entry.	R004152	One or more replica servers lack support for complex Modify DN operation refused.
R00)4139	Unable to locate effective ACLEntry for newSuperior entry.	R006001	LDAR Client ARI ani name has returned
R00)4140	Failure occurred while updating ACL and owner inheritance.	K000001	LDAP Client API api_name has returned an error code=return_code with an error message = error_string.
 R00)4141	Modify DN operation requires new rdn containing non-zero length string.	R006002	LDAP Client API api_name has returned an error code=return_code with an error message = error_string. Request will be retried.
R00)4142	Failure occurred while realigning DN attribute values to match renamed DNs with rc=return-code.	R006003	A decoding error has been encountered while decoding attibute(s): attr_type.
R00)4143	Could not realign DN attributes because no attribute types were found with DN syntax.	R006004	An encoding error has been encountered while encoding attibute(s): attr_type.
R00)4144	Retrieval of newSuperior object failed, despite apparent presence in backend.	R006005	A code page conversion error has been encountered.
R00)4145	The newSuperior DN is not permitted to be a referral object.	R006006	An unsupported control with OID=oid has been encountered.
R00)4146	Unable to complete DN realignment - generated SQL statement too large.	R006007	Critical extension with OID=oid not supported.
R00)4147	Cannot replicate complex Modify DN operation because a connection cannot be made with the replica. Modify DN operation is refused.	R006008	An LDAP Client API has returned an error code=return_code with an error message = error_string. Request may be retried.

R006009	The extended operation request with OID=oid requires the critical control	R007009	Error obtaining controls from ber.
	with OID=oid.	R007010	Control verify failed.
R006010	The extended operation request with OID=oid is not supported.	R007011	Error on dn_normalize of dn dn.
R006011	The extended operation request with OID=oid does not support the critical control with OID=oid.	R007012	Backend get_groups routine for be=backend failed.
	control with old-old.	R007013	GSSAPI bind crendentials are NULL.
R006012	The extended operation has encountered a cache miss.		Sasl bind should be in progress.
R006013	Mutex Exception. Errno=errno, Errno2=errno2, Error message = error_string.	R007015	Credentials do not match those in the configuration file.
R006014	Null Pointer Exception occurred.	R007016	No backend selected for entry dn.
 R006015	Reference Counting Exception	R007017	Backend bind failed for entry dn.
11000010	occurred.	R007018	CRAM-DM5 protocol error on bind.
R007000	Bind version version not supported.	R007019	DIGEST-MD5 protocol error on bind.
R007001	Only SIMPLE authentication method is supported for V2.	R007020	CRAM-DM5 bind with native authentication turned on is not supported.
R007002	Only EXTERNAL, GSSAPI, CRAM-MD5, and DIGEST-MD5 mechanisms are supported for SASL authentication methods.	R007021	DIGEST-MD5 bind with native authentication turned on is not supported.
R007003	Only SIMPLE and SASL authentication methods are supported for V3.	R007022	Error: Invalid CRAM-MD5 client response format.
R007004	Bind does not support manageDsalT control (managedsait).	R007023	Error: Could not locate uid attribute in schema to perform Mandatory Authentication bind.
R007005	Server not configured for client authentication.	R007024	Error: Found multiple entries with the
R007006	No client certificate dn available.		same uid value <i>uid</i> on Mandatory Authentication bind.
R007007	Server not configured for Kerberos GSSAPI bind.	R007025	Error: Authentication DN does not equal username DN in CRAM-MD5 bind.
R007008	Bind function not implemented for be=backend.		

R007026	Error: Mandatory Authentication bind not allowed to SDBM backend.
R007027	TLS is not supported on the connection.
R007028	TLS/SSL is active on the connection.
R007029	Other operations are outstanding for the connection.
R007030	Error: Multiple attribute-name attributes found in DIGEST-MD5 response.
R007031	Error: Required attribute attribute-name is missing in DIGEST-MD5 response.
R007032	Error: Quotations are missing around attribute attribute-name in DIGEST-MD5 response.
R007033	Error: Quotations are not needed around attribute attribute-name in DIGEST-MD5 response.
R007034	Error: Authorization DN does not equal username DN in DIGEST-MD5 bind.
R007035	Error: DIGEST-MD5 response length length is greater than 4096.
R007036	Error: DIGEST-MD5 response missing 'dn:' on 'authzid' attribute.
R007037	Error: DIGEST-MD5 response attribute attribute-name value is different than what was sent.
R007038	Credentials are not valid.
R007039	Referrals returned.

Part 3. Messages

Chapter 28. LDAP server messages

This part contains the messages returned by the LDAP server. The messages are ordered alphanumerically.

The new LDAP server messages for z/OS Release 4 are:

- GLD0207I GLD0240A
- GLD3133A GLD3143I
- GLD4013A GLD4020I
- GLD5006A

LDAP server messages (0000)

GLD0002l Configuration file successfully read.

Severity: Information

Explanation: The LDAP server successfully read the

configuration file.

System Action: The program continues.

Administrator Response: None.

GLD0004l Terminating slapd.

Severity: Information

Explanation: The LDAP server is ending, probably

due to a SIGTERM signal.

System Action: The program ends.

Operator Response: None.

Administrator Response: None.

GLD0005E The ber_scanf failed during operation:

operation.

Severity: Eventual Action

Explanation: The LDAP server is unable to process the requested operation because of a failure when

interpreting the request.

System Action: The program continues. The request

fails.

Administrator Response: Correct the request and try

again.

GLD0006E No values for type type.

Severity: Eventual Action

Explanation: The LDAP server is unable to process the requested operation because no values were

supplied for the specified type.

System Action: The program continues. The request

fails.

Administrator Response: Correct the request and try again.

GLD0008E Unrecognized database type (type).

Severity: Eventual Action

Explanation: The LDAP server encountered an error processing the configuration file. The specified database type is not supported.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail.

Administrator Response: Correct the configuration file and restart the server.

GLD0010I Reading configuration file file_name.

Severity: Information

Explanation: The LDAP server is processing the

specified configuration file.

System Action: The program continues.

Administrator Response: None.

GLD0011E command_name: line line_number:

option not valid or must appear inside a database definition. Ignored.

Severity: Eventual Action

Explanation: The LDAP server encountered an error processing the configuration file. The option on the specified line is either not recognized or is associated with a database definition and must appear in the database section of the configuration file. The specified line is ignored. If the line does not belong in a database section, verify that the option is spelled correctly in the configuration file. Misspelled options may not be recognized by the server.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration

parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and restart the server. Check the spelling of the configuration option on the specified line in the configuration file. Also verify that this option appears in the correct section of the configuration file.

GLD0012E command name: line line number: incorrect configuration line. Ignored.

Severity: Eventual Action

Explanation: The LDAP server encountered an error processing the configuration file because the specified line is not correct. The line is ignored.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and restart the server.

GLD0013E

command_name: line line_number: incorrect number of parameters specified.

Severity: Eventual Action

Explanation: The LDAP server encountered an error processing the configuration file because the specified line does not supply all of the required parameters.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and restart the server.

GLD0014A Unable to open file file_name. Try specifying the full path name.

Severity: Immediate Action

Explanation: The LDAP server or a LDAP utility is unable to open the specified configuration file.

System Action: The program ends.

Operator Response: Correct the file name and restart the server, or contact the administrator.

Administrator Response: Determine the reason LDAP was unable to open the file. Possible reasons include misspelled file name, path to the file not fully specified, or file permissions incorrectly set. Correct the problem and restart the server.

GLD0015E

command_name: line line_number: unknown directive directive outside database definition. Ignored.

Severity: Eventual Action

Explanation: The LDAP server encountered an error processing the configuration file because the specified line contains an unrecognized directive or keyword. The specified line is ignored.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and restart the server.

GLD0016A A ber_alloc failed.

Severity: Immediate Action

Explanation: The LDAP server is unable to allocate the necessary storage to continue processing the request.

System Action: The program ends.

Operator Response: Increase the storage for the LDAP server and restart the server. If the problem persists, contact the service representative.

GLD0019A Error while trying to allocate memory.

Severity: Immediate Action

Explanation: The LDAP server is unable to allocate the necessary storage to continue processing.

System Action: The program ends.

Operator Response: Increase the storage for the LDAP server and restart the server. If the problem persists, contact the service representative.

GLD0020A

Exceeded maximum number of connections, currently set at

max_connections.

Severity: Immediate Action

Explanation: The LDAP server is unable to process

the request because all available connections are in

System Action: The program continues. The request

Administrator Response: Submit the request again.

GLD0021A Unable to create the necessary threads for threadpool.

Severity: Immediate Action

Explanation: The LDAP server is unable to obtain the necessary resources to create the threads in this threadpool. When this failure occurs during the startup of the LDAP Server, the program will end.

System Action: The program ends.

Operator Response: Verify the OMVS Kernel is operating correctly. Save the dump and contact the system programmer.

Administrator Response: If this error occurs with the commThreadPool, lower the number of commThreads in the configuration file or increase other tunable OMVS parameters such as MAXTHREADTASKS or MAXTHREADS. If this error occurs with the pcThreadPool, lower the number of pcThreads in the configuration file or increase other tunable OMVS parameters such as MAXTHREADTASKS or MAXTHREADS. If this error occurs with the replThreadPool, increase tunable OMVS parameters such as MAXTHREADTASKS or MAXTHREADS. If this problem persists, contact the service representative.

GLD0022I LDAP_server_version Starting slapd.

Severity: Information

Explanation: The LDAP server is starting. **System Action:** The program continues.

Operator Response: None.

Administrator Response: None.

GLD0027A slapd unable to start because all backends failed to configure.

Severity: Immediate Action

Explanation: The LDAP server is unable to start because the defined backends have not configured successfully.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Check for other messages regarding errors during configuration. Correct the configuration file and restart the server.

GLD0028E Configuration error: server using port port number for both SSL and non-SSL.

Severity: Eventual Action

Explanation: The LDAP server is listening on the specified port for both secure and nonsecure requests.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the port values specified in the configuration file and restart the server.

GLD0038A A command that is not supported or is not valid was entered from the console.

Severity: Immediate Action

Explanation: The LDAP server received a command from the operator console that is not supported or is not

System Action: The program continues.

Operator Response: Correct the command and try again.

GLD0039A The __console() function failed with errno errno.

Severity: Immediate Action

Explanation: The LDAP server received an unexpected return code from system function console(). The z/OS C/C++ Run-Time Library Reference, SC28-1663 contains more information about the **errno**. Error messages normally written to the console will only appear in the log.

System Action: The program continues.

Operator Response: Save the dump, if any, and contact the system programmer. If the problem persists, contact the service representative.

GLD0041A An administrator DN must be specified in the adminDn line.

Severity: Immediate Action

Explanation: The LDAP server encountered a blank adminDN parameter in the configuration file.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the adminDN parameter in the configuration file and restart the server. GLD0043A Unable to connect to replica

replica name on port port number. Please verify that the replica is started.

Severity: Immediate Action

Explanation: The LDAP server was unable to connect to the specified server to perform replication.

System Action: The program continues. Replication to the specified server cannot continue.

Operator Response: Verify that the replica server is started.

Administrator Response: Verify that the replica server is started, or contact the operator to start the replica server. Verify that the replica server information is correct.

GLD0044E

Error error_value occurred during replication to replica name: add failed on entry distinguished_name.

Severity: Eventual Action

Explanation: The LDAP server was unable to replicate the specified add operation to the replica server. The replication attempt will be tried again.

System Action: The program continues.

Administrator Response: Verify that the replica server information is correct. If the failure continues, contact the service representative.

GLD0045E

Error error value occurred during replication to replica_name: delete failed on entry distinguished_name.

Severity: Eventual Action

Explanation: The LDAP server was unable to replicate the specified delete operation to the replica server. The replication attempt will be tried again.

System Action: The program continues.

Administrator Response: Verify that the replica server information is correct. If the failure continues, contact the service representative.

GLD0046E

Error error_value occurred during replication to replica_name: modrdn failed on entry distinguished_name.

Severity: Eventual Action

Explanation: The LDAP server was unable to replicate the specified modify RDN operation to the replica server. The replication attempt will be tried again.

System Action: The program continues.

Administrator Response: Verify that the replica server information is correct. If the failure continues, contact the service representative.

GLD0047E

Error error_value occurred during replication to replica_name: modify failed on entry distinguished_name.

Severity: Eventual Action

Explanation: The LDAP server was unable to replicate the specified modify operation to the replica server. The replication attempt will be tried again.

System Action: The program continues.

Administrator Response: Verify that the replica server information is correct. If the failure continues, contact the service representative.

GLD0048A

Creation of socket failed; errno errno (errno_string).

Severity: Immediate Action

Explanation: The LDAP server received the specified error from system function **socket()**. Refer to the *z/OS* C/C++ Run-Time Library Reference, SC28-1663 for an explanation of the errno returned.

System Action: Initialization continues for all other requested communication interfaces. If no communication interface is successfully initialized, the program ends.

Operator Response: Ensure TCP/IP is operating correctly. Save the diagnostic information and contact the system programmer.

Administrator Response: If the problem persists, contact the service representative.

GLD0049A

Attempt to setsockopt(socket option) failed; errno errno (errno_string).

Severity: Immediate Action

Explanation: The LDAP server received the specified error from system function setsockopt(). Refer to the z/OS C/C++ Run-Time Library Reference, SC28-1663 for an explanation of the errno returned.

System Action: The program continues.

Operator Response: Ensure TCP/IP is operating correctly. Save the diagnostic information and contact the system programmer.

Administrator Response: If the problem persists, contact the service representative.

GLD0050A

Attempt to bind failed; errno errno (errno_string).

Severity: Immediate Action

Explanation: The LDAP server received an error from system function bind(). Refer to the z/OS C/C++

Run-Time Library Reference, SC28-1663 for an explanation of the **errno** returned.

System Action: Initialization continues for all other requested communication interfaces. If no communication interface is successfully initialized, the program ends.

Operator Response: Ensure TCP/IP is operating correctly. Save the diagnostic information and contact the system programmer.

Administrator Response: If the problem persists, contact the service representative.

GLD0051A The listen() failed; errno errno (errno_string).

Severity: Immediate Action

Explanation: The LDAP server received an error from system function **listen()**. Refer to the *z/OS C/C++ Run-Time Library Reference*, SC28-1663 for an explanation of the **errno** returned.

System Action: Initialization continues for all other requested communication interfaces. If no communication interface is successfully initialized, the program ends.

Operator Response: Ensure TCP/IP is operating correctly. Save the diagnostic information and contact the system programmer.

Administrator Response: If the problem persists, contact the service representative.

GLD0052I Configuration read securePort port_number.

Severity: Information

Explanation: The LDAP server has assigned the securePort to the specified value based on the value read from the configuration file.

System Action: The program continues.

Administrator Response: None.

GLD0053I Configuration read security of

security_value.

Severity: Information

Explanation: The LDAP server has assigned the security to the specified value based on the value read from the configuration file.

System Action: The program continues.

Administrator Response: None.

GLD00601 Non-SSL port override from command line. value = value.

Severity: Information

Explanation: The LDAP server has assigned the nonsecure port to the specified value based on the value supplied on the command line.

System Action: The program continues.

Administrator Response: None.

GLD0061I SSL port override from command line, value = value.

Severity: Information

Explanation: The LDAP server has assigned the secure port to the specified value based on the value supplied on the command line.

System Action: The program continues.

Administrator Response: None.

GLD0063A Error error_code reported by SSL initialization.

Severity: Immediate Action

Explanation: The LDAP server was unable to complete initialization required for secure communications.

System Action: Communication interfaces that require SSL will not be activated. Communication interfaces that support both secure and non-secure communication will be activated for non-secure communication only. If no interfaces are activated then the program will end.

Operator Response: The error code is documeted in the IdapssI.h file shown in the LDAP Client Programming manual. If the error description indicates a correctable error then correct it and restart the server. Otherwise, contact the service representative.

GLD0064A SSL detected an error reading file_name.

Severity: Immediate Action

Explanation: The LDAP server was unable to read the key database file or key ring required for secure communications. Secure communications cannot continue.

System Action: Communication interfaces that require SSL will not be activated. Communication interfaces that support both secure and non-secure communication will be activated for non-secure communication only. If no interfaces are activated then the program will end.

Operator Response: Verify that the file system is operating correctly.

Administrator Response: Ensure that the file

permissions on the SSL keyring file are correct.

GLD0065A SSL encountered an error opening file name.

Severity: Immediate Action

Explanation: The LDAP server was unable to open the key database file or key ring required for secure communications. Secure communications cannot continue

System Action: Communication interfaces that require SSL will not be activated. Communication interfaces that support both secure and non-secure communication will be activated for non-secure communication only. If no interfaces are activated then the program will end.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Ensure that the file permissions on the SSL key database file are correct and that it exists.

GLD0066A SSL key database file file_name is in an unknown format.

Severity: Immediate Action

Explanation: The System SSL runtime library is unable to decrypt a key database entry. Either the supplied database password is incorrect or the database is damaged. Secure communications cannot continue.

System Action: Communication interfaces that require SSL will not be activated. Communication interfaces that support both secure and non-secure communication will be activated for non-secure communication only. If no interfaces are activated then the program will end.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Ensure that the correct key database password is specified in the configuration file. Recreate the database if the error persists.

GLD0067A The password supplied for SSL key database file file_name is not correct.

Severity: Immediate Action

Explanation: The LDAP server found that the password supplied for the key database file is incorrect. Secure communications cannot continue.

System Action: Communication interfaces that require SSL will not be activated. Communication interfaces that support both secure and non-secure communication will be activated for non-secure communication only. If no interfaces are activated then the program will end.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the password for the key database file in the configuration file or the password stash file and restart the server.

GLD0069A Error error code reported attempting SSL handshake.

Severity: Immediate Action

Explanation: The LDAP server was unable to complete initialization of the socket required for secure communications.

System Action: The program continues.

Operator Response: The error code is documented in the Idapssl.h file shown in the LDAP Client Programming manual. If the error description indicates a client error then correct it and restart the client. If the error description indicates a correctable server error then correct it and restart the server. Otherwise, contact the service representative.

GLD0070A An incorrect SSL certificate label label was specified.

Severity: Immediate Action

Explanation: The LDAP server found that the certificate label supplied for the key database file or key ring is incorrect. Secure communications cannot continue.

System Action: The program continues. Any communication inferface related to SSL will not continue

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the certificate label in the configuration file and restart the server.

GLD0072A No SSL ciphers matched both the client and server cipher specifications.

Severity: Immediate Action

Explanation: The client and server cipher specifications do not contain at least one value in common. This error can also occur if no SSL protocols are enabled or if all of the enabled protocols have empty cipher specifications.

System Action: The program continues. Secure communication between the server and the client will

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Ensure that the client and the server have at least one cipher specification in common. If the server needs additional ciphers then update the cipher specification in the configuration file and restart the server. If the client needs additional

ciphers then correct that condition and retry the client.

GLD0073A The default SSL certificate has expired in file name.

Severity: Immediate Action

Explanation: The LDAP server found that the default certificate in the key database file or key ring is no longer valid.

System Action: The program continues. Any communication inferface related to SSL will not continue.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Refresh the certificate and restart the server.

GLD0075A The default SSL certificate is of an unsupported type in file_name.

Severity: Immediate Action

Explanation: The LDAP server found that the default certificate in the key database file or key ring is not the correct type.

System Action: The program continues. Any communication inferface related to SSL will not continue.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Obtain a supported certificate and restart the server.

GLD0076A No SSL certificate exists in file name.

Severity: Immediate Action

Explanation: The LDAP server found that no certificate exists in the key database file or key ring.

System Action: The program continues. Any communication inferface related to SSL will not continue.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Refresh the certificate and restart the server.

GLD0077A The underlying socket was closed.

Severity: Immediate Action

Explanation: The LDAP server was unable to complete secure communications because the socket is closed.

System Action: The program continues.

Operator Response: None.

Administrator Response: Contact the service representative.

GLD0083I Successfully reconnected to replica host replica name on port port number.

Severity: Information

Explanation: The LDAP server was able to restart replication to the specified replica server.

System Action: The program continues.

Administrator Response: None.

GLD0084A Connection to replica replica_name on port port_number has failed. Verify that the replica is started.

Severity: Immediate Action

Explanation: The LDAP server detected that the connection to the specified replication server ended.

System Action: The program continues. Replication to the specified server cannot continue.

Operator Response: Verify that the replica server is started.

Administrator Response: Verify that the replica server is started or contact the operator to start the replica server. Verify that the replica server information is correct.

GLD0089E Attention: configuration file file_name is empty.

Severity: Eventual Action

Explanation: The LDAP server encountered an empty configuration file.

System Action: The program ends if the empty file is the configuration file. The program continues if the empty file is an included configuration file.

Administrator Response: Ensure the correct configuration files are specified.

GLD0090A A modify command that failed was entered from the console: command.

Severity: Immediate Action

Explanation: The LDAP server received a modify command from the operator, but the command failed. Either the syntax was wrong or the processing invoked by the command experienced a problem.

System Action: The program continues.

Operator Response: Correct the syntax or determine why the processing of the command failed.

GLD00911 Successfully set debug level to

debug_level from console command.

Severity: Information

Explanation: The LDAP server set the debug level to

the specified value.

System Action: The program continues.

Operator Response: None.

Administrator Response: None.

GLD0092A

Unable to open any configuration file. No configuration file specified at startup, tried DDname DD_name and default name file_name.

Severity: Immediate Action

Explanation: The LDAP program is unable to open any configuration file. No configuration file name was specified at startup using the -f option. The program tried the specified DDname and then the default configuration file, but was unable to find any configuration file. The program cannot start without a configuration file.

System Action: The program ends.

Operator Response: Correct the configuration file name and restart the server or contact the administrator.

Administrator Response: Correct the file name or permissions and restart the server.

GLD0107A The configuration file produces a loop within the include statements.

Severity: Immediate Action

Explanation: A loop was detected during the processing of the included configuration files during server startup.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Check the configuration file

include statements.

GLD0108E No object class was specified for entry entry dn.

Severity: Eventual Action

Explanation: All entries must specify an object class.

System Action: The program continues. The request

fails.

Administrator Response: Verify the 1dif file syntax.

GLD0109E The required attribute attribute name is missing for entry entry_dn.

Severity: Eventual Action

Explanation: An attribute required by an object class specified in this entry was not provided.

System Action: The program continues. The request

fails.

Administrator Response: Verify the 1dif file syntax. Refer to the schema files for a list of the attributes required for each object class.

GLD0110E The attribute attribute name is not allowed for entry entry_dn.

Severity: Eventual Action

Explanation: The specified attribute is not allowed by the object classes specified in this entry.

System Action: The program continues. The request

Administrator Response: Verify the 1dif file syntax. Refer to the schema files for a list of the attributes allowed for each object class.

GLD0111E

An error occurred while processing schema file file_name: line line_number syntax :oc clause ::= objectclass ocname [requires attrlist] [allows attrlist]

Severity: Eventual Action

Explanation: The LDAP server object class schema file contains a schema definition that is not valid.

System Action: The program ends.

Administrator Response: Correct the object class definition in the specified schema file and restart the server.

GLD0114I A client sent nonsecured communications to the SSL port.

Severity: Information

Explanation: The LDAP server determined that a client sent unencrypted data to the secure port. The request from the client is ended.

System Action: The program continues. The request fails.

Administrator Response: If a command utility such as Idapsearch is called by the client and secure communications is intended, make sure the -Z (use secure communications) parameter is specified. If the client does not intend to use secure communications, then specify the nonsecure port. If secure communications is intended, make sure the client calls **Idap_ssI_start**. If secure communications is not intended then specify the nonsecure port.

GLD0115I

Workload Manager enablement initialization successful for group=sysplex_groupname, server=sysplex_servername on host host_name.

Severity: Information

Explanation: The LDAP server successfully registered with Sysplex Workload Manager. The LDAP server will operate in multi-server mode. Multiple concurrent servers may be running on a given OS/390 image, and multiple concurrent servers may be running on multiple OS/390 images in a parallel sysplex, all of which use the same LDAP DB2 database. Note that no replication may be performed by any of these servers. If replication is desired, only one server instance may be running which uses a given LDAP DB2 database, and the server must be configured to run in single-server mode.

System Action: The program continues.

Operator Response: None.

Administrator Response: None.

GLD0116E

Workload Manager enablement initialization failed for group=sysplex_groupname, server=sysplex_servername on host host_name. No Workload Manager support will be activated for this server. RC = return_code.

Severity: Eventual Action

Explanation: The LDAP server registration with Sysplex Workload Manager failed.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Contact the service representative and provide the return code in this message.

GLD0117I

Workload Manager enablement termination successful for group=sysplex_groupname, server=sysplex_servername on host host_name.

Severity: Information

Explanation: The LDAP server successfully unregistered from Sysplex Workload Manager.

System Action: The program continues.

Administrator Response: None.

GLD0118E

Workload Manager enablement termination failed for group=sysplex_groupname, server=sysplex_servername on host host_name. RC = return_code.

Severity: Eventual Action

Explanation: The LDAP server was unable to unregister from Sysplex Workload Manager.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: If the problem persists,

contact the service representative.

GLD0119E

Workload Manager enablement initialization failed for group=sysplex_groupname, server=sysplex_servername. No Workload Manager support will be activated for this server. RC = return_code.

Severity: Eventual Action

Explanation: The LDAP server was unable to register with Sysplex Workload Manager. This is an internal program error. Contact the service representative with the RC value displayed.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Contact the service representative.

GLD0120E

Workload Manager enablement termination failed for group=sysplex_groupname, server=sysplex_servername. RC = return_code.

Severity: Eventual Action

Explanation: The LDAP server was unable to deregister from Sysplex Workload Manager. This is an internal program error. Contact the service representative with the RC value displayed.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator

or see the Administrator Response.

Administrator Response: Contact the service

representative.

GLD0121E The object class object class specified

for entry distinguished name is not defined in the schema.

Severity: Eventual Action

Explanation: An object class specified in a request for the given DN is not defined in the current schema.

System Action: The program continues. The request

fails.

Administrator Response: Verify the object class in the request and the object classes defined in the schema configuration files and try the request again.

GLD0122I Slapd is ready for requests.

Severity: Information

Explanation: The LDAP server is listening and is

ready for requests.

System Action: The program continues.

Operator Response: None.

Administrator Response: None.

GLD0123E

Workload Manager enablement initialization failed for group=sysplex_groupname, server=sysplex_servername on host host_name because this group/server combination is already registered on this host. No Workload Manager support will be activated for this server. RC = return_code.

Severity: Eventual Action

Explanation: The LDAP server registration with Sysplex Workload Manager failed because another server has already been registered by the same name in the same sysplex group.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator

or see the Administrator Response.

Administrator Response: Ensure that the sysplexServerName in the server configuration file (slapd.conf) is unique among all servers started with the same sysplexGroupName. If the problem persists, contact the service representative.

GLD0124I

Dynamic workload management enabled. Server will operate in multi-server mode. sysplexServerName

= sysplex server name, sysplexGroupName = sysplex_group_name.

Severity: Information

Explanation: The LDAP server will operate in multi-server mode. Multiple concurrent servers may be running on a given z/OS image, and multiple concurrent servers may be running on multiple z/OS images in a parallel sysplex, all of which use the same LDAP DB2 database. Note that no replication may be performed by any of these servers. If replication is desired, only one server instance may be running which uses a given LDAP DB2 database, and the server must be configured to run in single-server mode.

System Action: The program continues.

Administrator Response: None.

GLD0125E

Only one of the sysplexGroupName or sysplexServerName keywords is present in the configuration file.

Severity: Eventual Action

Explanation: One of the sysplex keywords sysplexGroupName or sysplexServerName was found in the server configuration file. If either of these keywords is present, the other must also be present, and both keywords must be accompanied by non-null arguments.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the configuration file and restart the server.

GLD0126A

The length of sysplex_keyword must be less than or equal to maximum_argument_length.

Severity: Immediate Action

Explanation: The length of the sysplexGroupName argument must not be greater than 18 characters. The length of the sysplexServerName argument must not be greater than 8 characters.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the configuration file and restart the server.

GLD0127A

Workload Manager enablement initialization failed due to memory allocation error. No Workload Manager support will be activated for this server.

Severity: Immediate Action

Explanation: The LDAP server could not allocate memory needed to register with Workload Manager. The LDAP server will continue to operate, but no Workload Manager functions will be available.

System Action: The program continues.

Operator Response: Increase the storage for the LDAP server and restart the server. If the problem persists, contact the service representative.

GLD0128A

The SSL certificate sent by the client or the server certificate in file name is not valid.

Severity: Immediate Action

Explanation: The LDAP server found that the certificate sent by the client or the certificate in the key database file or key ring is not valid.

System Action: The program continues. The request fails.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Verify that the server certificate in the key database file or key ring is still valid. Things to look for in the server certificate information are the issuer Certificate Authority and the expiration time of the certificate.

GLD0129E

The value for the 'maxConnections' option is out of range (min_Connection, max_Connection). The default (max_Connection) will be used.

Severity: Eventual Action

Explanation: The LDAP server is unable to process the request because the value of maxConnections is out of range. The server will continue with the default value

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0131E

The value for the sizelimit option is not numeric. The default (default_sizelimit) will be used.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the sizelimit option of the

configuration file is not numeric. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0132E

The value for the timelimit option is not numeric. The default (default_timelimit) will be used.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the timelimit option of the configuration file is not numeric. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0135E

The value specified for the readonly option is not valid. The default (default_readonly) will be used.

Severity: Eventual Action

Explanation: The LDAP server encountered an error during processing of the readonly option of the configuration file. The default value will be used.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0136A

The value specified for the masterserverDN option is not valid.

Severity: Immediate Action

Explanation: The LDAP server encountered an error during processing of the masterserverDN option of the configuration file. The masterserverDN cannot have a NULL value.

System Action: The program continues. However, configuration may fail. Replication will not be configured.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the configuration file and restart the server.

GLD0137A

The value specified for the masterserverDN option is 'cn=Anybody'. This is not permitted.

Severity: Immediate Action

Explanation: The LDAP server encountered an error during processing of the masterserverDN option of the

configuration file. The masterserverDN cannot be 'cn=Anybody'.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the configuration

file and restart the server.

GLD0138E The value for the maxConnections option is not numeric. The default

(default_max_connections) will be used.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the maxConnections option of the configuration file is not numeric. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0141E A backend (backend_address) of type backend_type failed to configure.

Severity: Eventual Action

Explanation: The LDAP server encountered an error during configuration of the specified backend. This backend will not be available when the server starts.

System Action: The program continues. Other backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Check for other messages regarding errors during configuration. Correct the configuration file and restart the server.

GLD0142E The value URL_value for the

configuration_option option is not a valid

URL.

Severity: Eventual Action

Explanation: The LDAP server determined that the value provided for the specified option of the configuration file is not valid. An LDAP URL is expected.

System Action: The program continues. However, configuration may fail.

Operator Response: None.

Administrator Response: Correct the configuration

file parameter and restart the server.

GLD0143A The LDAP program cannot create required configuration structures.

Severity: Immediate Action

Explanation: The LDAP program encountered an error when establishing a required configuration

structure. The program cannot start.

System Action: The program ends.

Operator Response: Increase the storage for the LDAP server and restart the server. If the problem persists, contact the service representative.

Administrator Response: Ensure adequate memory for the program. Correct any other reported errors and restart the program.

GLD0144A The LDAP server encountered an error during configuration.

Severity: Immediate Action

Explanation: The LDAP server encountered an error during configuration. The program cannot start. See other messages regarding errors.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response. If the problem persists, contact the service representative.

Administrator Response: Examine other messages and correct any other reported errors. Restart the LDAP server.

GLD0146E The value for the -s command line option is not numeric. Value is ignored.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the -s option on the command line invocation of the program is not numeric. The value is ignored. The -s option identifies the secure port.

System Action: The program continues.

Administrator Response: Correct the command line option and restart the server.

GLD0147E The value for the -p command line option is not numeric. Value is ignored.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the -p option on the command line invocation of the program is not numeric. The value is ignored. The -p option identifies the port.

System Action: The program continues.

Administrator Response: Correct the command line option and restart the server.

GLD0148A

The dilload function failed loading path loadpath with errno=errno, error string=error string.

Severity: Immediate Action

Explanation: The LDAP server was unable to load the requested library. Refer to the z/OS C/C++ Run-Time Library Reference, SC28-1663 for explanation of the errno.

System Action: The program continues. However, configuration may fail.

Operator Response: Contact the LDAP Administrator or see the Administrator Response. If the problem persists, contact the service representative.

Administrator Response: Correct the error and restart the server.

GLD0149E

The value specified for the extendedGroupSearching option is not valid. The default value of (default_extendedGroupSearching) will be used.

Severity: Eventual Action

Explanation: The LDAP server encountered an error during processing of the extendedGroupSearching option of the configuration file. The default value will be used.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0150A The LDAP program requires a TDBM backend but none is configured.

Severity: Immediate Action

Explanation: A TDBM backend must be configured for this LDAP program to run. Either no TDBM backend was configured or an error was encountered during TDBM configuration. The LDAP program cannot start.

System Action: The program ends.

Operator Response: None.

Administrator Response: Ensure TDBM configuration is present. Correct any other reported errors and restart the program.

GLD0151A

Object class object_class requires attribute type attribute_type which is not defined.

Severity: Immediate Action

Explanation: The LDAP server encountered an error during processing of the object class definitions.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Correct the schema definition by correcting or removing the attribute type from the object class definition or adding the attribute type to the schema and restart the server.

GLD0154E **Error code** *error_code* **from odbc string**: "odbc_string" substring.

Severity: Eventual Action

Explanation: An error occurred when performing DB2 operations. The request fails. There may be an additional message with information about SQL return codes.

System Action: The program continues. The request fails.

Administrator Response: Evaluate and resolve the DB2 problem with the information provided. If the problem cannot be resolved, contact the service representative.

ODBC error, SQL data is: native return GLD0155E code=SQL_code, SQL state=SQL_state, **SQL** message=SQL_message.

Severity: Eventual Action

Explanation: An error occurred when performing DB2 operations. The information in the message is the data available from SQL at the time of the error.

System Action: The program continues. The request fails.

Administrator Response: Evaluate and resolve the DB2 problem with the information provided. If the problem cannot be resolved, contact the service representative.

GLD0156A Unable to normalize administrator DN: admin dn.

Severity: Immediate Action

Explanation: The LDAP program encountered an error while attempting to normalize the administrator DN from the configuration file. A possible reason for the error is no equal sign in the relative DN. admin dn will be binary if unable to convert the string.

System Action: The program continues. However, configuration may fail.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Determine the reason for the error. Correct the adminDN in the configuration file and restart the server.

GLD0157A **Unable to normalize master server DN:**

master server dn.

Severity: Immediate Action

Explanation: The LDAP program encountered an error while attempting to normalize the master server DN from the configuration file. A possible reason for the error is no equal sign in the relative DN. admin_dn will be binary if unable to convert the string.

System Action: The program continues. However, configuration may fail.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Determine the reason for the error. Correct the masterServerDN in the configuration file and restart the server.

GLD0158A Unable to normalize suffix: suffix.

Severity: Immediate Action

Explanation: The LDAP program encountered an error while attempting to normalize a suffix from the configuration file. A possible reason for the error is no equal sign in the relative DN. admin_dn will be binary if unable to convert the string.

System Action: The program continues. However, configuration may fail.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Determine the reason for the error. Correct the suffix in the configuration file and restart the server.

GLD0159A Unknown exception caught during configuration.

Severity: Immediate Action

Explanation: The LDAP program encountered an unknown error during configuration.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: It may be possible to determine the error using LDAP debug tracing. If unable to resolve the problem using the debug trace, the service representative should be contacted.

GLD0160E

A replica with dn: replica dn, host name: hostname and port: port duplicates existing replica with dn: replica dn. Existing replica will be used.

Severity: Eventual Action

Explanation: The identified replica object has the same host name and port as an existing replica object, and the bind dn and credentials also match those in the existing replica object. The existing replica object will be used to control replication. While the new replica will be added to the directory, it will not be used to control replication.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0161E

A replica with dn: replica_dn, host name: hostname and port: port attempted to duplicate existing replica with dn: replica_dn. New replica not added to directory.

Severity: Eventual Action

Explanation: The identified replica object has the same host name and port as an existing replica object. One of the following reasons applies:

- The new replica object has the same host and port as another replica object in the directory, but the bind dn and credentials do not match those in the existing replica object.
- The new replica object duplicates another replica object that was recently deleted from the directory. Clean up processing for the recently removed replica object has not completed.

The existing replica object remains active.

System Action: The program continues. However, the duplicate replica definition is not added to the directory.

Administrator Response: A replica with the same host name and port as an existing replica can only be added if the bind dn and credentials in the two entries also match. Correct the replicaObject entry being added and try the add again. If a replica object has been recently deleted, wait ten minutes to allow clean up processing to complete and try the request again.

GLD0162I

Backend of type: backend_type serving suffix: suffix_dn does not support reporting of backend capabilities.

Severity: Information

Explanation: The LDAP server has loaded a backend which does not contain the necessary programming support to report the capabilities of that backend.

System Action: The program continues.

Administrator Response: None.

GLD01631 Backend capability listing follows:

Severity: Information

Explanation: The LDAP server has completed loading backends and will report the capabilities of those backends which have implemented the necessary programming support to do so.

System Action: The program continues.

Administrator Response: None.

GLD0164I Backend capability listing ended.

Severity: Information

Explanation: The LDAP server has completed reporting the capabilities of those backends which have implemented the necessary programming support to do

System Action: The program continues.

Administrator Response: None.

GLD0165I Capability: attribute_type Value:

attribute value

Severity: Information

Explanation: The LDAP server is reporting the capabilities of backends which have implemented the necessary programming support to do such reporting. The Capability indicates what capability attribute is being reported for the backend, and the Value expresses the attribute value for that particular capability. Following is a partial list of common capabilities which may be reported, accompanied by a brief description of each:

- Backend ID a string used to uniquely identify the backend.
- · Build Date and Time the date on which the backend was compiled, in the format yyyy-mm-ddhh.mm.ss.000000.
- Apar Level if the backend was compiled as part of an APAR, the APAR number is reported by the backend.
- · Release a release identifier associated with the backend.
- · Version a version identifier associated with the backend.
- · Dialect an indicator of the internal interface version to which the backend is programmed.
- · Berdecoding indicates whether the server front-end code should decode the incoming BER value in string format or in binary format before passing the values to the backend.

- ExtendedGroupSearching indicates whether the backend is capable of participation in extended group membership searching on a client bind request.
- SupportedControls lists which controls are supported by the backend.
- SupportedExtension lists which extendedRequests and extendedResponses are supported by the backend.

System Action: The program continues.

Administrator Response: None.

GLD0166I Backend type: backend_type, Backend

ID: backend identifier

Severity: Information

Explanation: The LDAP server is reporting the capabilities of backends. Backend type is the type of backend configured in the server configuration file (that is, rdbm, sdbm, tdbm, and so on). Backend ID is the identity of the backend.

System Action: The program continues.

Administrator Response: None.

GLD0167I End of capability listing for Backend

type: backend_type, Backend ID:

backend_identifier.

Severity: Information

Explanation: The LDAP server has completed reporting the capabilities of the current backend. Backend type is the type of backend configured in the server configuration file (that is, rdbm, sdbm, tdbm, and so on). Backend ID is the identity of the backend.

System Action: The program continues.

Administrator Response: None.

GLD0168I

Backend type backend_type: Backend suffix: backend suffix lacking required capability attribute(s) or value(s) and will be unloaded.

Severity: Information

Explanation: The LDAP server detected a backend which implemented the necessary programming support to report capabilities of the backend, and which is lacking one or more required capability attributes or attribute values. This backend will be unloaded and will be unavailable to the server.

System Action: The program continues.

Administrator Response: Under normal operation, this error should not occur. Contact the service representative.

GLD0169A Acquisition or check of backend capabilities failed for backend of type:

backend type.

Severity: Immediate Action

Explanation: During its check of backend capabilities, the LDAP server encountered an unrecoverable error.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: If the problem persists for the same Backend Type remove the database stanza for this Backend Type from the server configuration file to permit starting the server without this backend (if desired), and contact the service representative.

GLD0170I Kerberos authentication support has been enabled.

Severity: Information

Explanation: The LDAP server will be configured for

Kerberos authentication.

System Action: The program continues.

Administrator Response: None.

GLD01711 Kerberos authentication support has not been enabled.

Severity: Information

Explanation: The LDAP server will not be configured

for Kerberos authentication.

System Action: The program continues.

Administrator Response: None.

GLD0173E Server was unable to acquire Kerberos credentials.

Severity: Eventual Action

Explanation: Kerberos credentials could not be obtained for the server. Kerberos support will be

disabled.

System Action: The program continues.

Operator Response: Make sure the Kerberos KDC

has been started.

Administrator Response: If not using Kerberos, remove Kerberos information from the configuration file and restart the server.

GLD0174E Server is unable to acquire Kerberos credentials without a principal.

Severity: Eventual Action

Explanation: The server's Kerberos principal name serverKrbPrinc was not specified in the configuration file. Kerberos support will be disabled.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Add the server's Kerberos principal to the configuration file and restart the server.

GLD0175E The value for the 'commThreads' option is not numeric. The default

(default_comm_threads) will be used.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the commThreads option of the configuration file is not numeric. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0176E The value for the 'commThreads'

option is out of range (min_threads, max_threads). The default

(default comm threads) will be used.

Severity: Eventual Action

Explanation: The LDAP server encountered an error when processing the commThreads parameter of the configuration file. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0177E The value for the

'idleConnectionTimeout' option is not numeric. The default of indefinite (0) will be used.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the idleConnectionTimeout option of the configuration file is not numeric. The server will continue with the default value.

System Action: The program continues. Connections will not timeout.

Administrator Response: Correct the configuration file and restart the server.

GLD0178E The value for the

'idleConnectionTimeout' option is less than min_timeout. The default of indefinite (0) will be used.

Severity: Eventual Action

Explanation: The LDAP server encountered an error when processing the idleConnectionTimeout parameter of the configuration file. The server will continue with the default value.

System Action: The program continues. Connections will not timeout.

Administrator Response: Correct the configuration file and restart the server.

GLD0179A An internal error was detected in the communication area start-up table for tower 'tower_ID'.

Severity: Immediate Action

Explanation: The LDAP server detected an internal error in the communication area start-up table.

System Action: The program continues.

Operator Response: Save the diagnostic information and contact the administrator.

Administrator Response: Contact the service representative.

GLD0181E The 'threadParm' parameter is no longer supported.

Severity: Eventual Action

Explanation: This thread parameter is now obsolete and is being ignored. Use the commThreads parameter to set the number of threads to be used for communications between the clients and the backend(s).

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0182A slapd was unable to start because of communication errors.

Severity: Immediate Action

Explanation: The LDAP server was unable to start any of the interfaces used for communication with clients.

System Action: The program ends.

Operator Response: Save the diagnostic information and contact the administrator.

Administrator Response: Check for other messages regarding errors during configuration. Correct the

configuration file and restart the server.

GLD0183A The dilload function failed loading DLL dll_name; errno not available.

Severity: Immediate Action

Explanation: The LDAP server was unable to load the requested library. Due to the order of processing, errno information is not available. The most common reason for the error is that the specified DLL cannot be found in the libraries being searched.

System Action: The program continues. However, support provided by the DLL will not be available.

Operator Response: Contact the LDAP Administrator or see the Administrator Response. If the problem persists, contact the service representative.

Administrator Response: Correct the error and restart the server.

GLD0184I Connections allowed only on the secure port.

Severity: Information

Explanation: The LDAP server is allowing socket

communications on the secure port only.

System Action: The program continues.

Administrator Response: None.

GLD0185I Connections allowed only on the nonsecure port.

Severity: Information

Explanation: The LDAP server is allowing socket

communications on the nonsecure port only.

System Action: The program continues.

Administrator Response: None.

GLD0186I Connections allowed on both the secure and nonsecure port.

Severity: Information

Explanation: The LDAP server is allowing socket communications on both the secure and nonsecure

ports.

System Action: The program continues.

Administrator Response: None.

GLD01871 Configuration parameter

> configParameter is being ignored since it was specified with the listen parameter.

Severity: Information

Explanation: When a valid listen parameter is specified in the configuration file, it takes precedence over what has been specified for the Configuration parameter. The LDAP server will configure according to what was been specified for the listen parameter.

System Action: The program continues.

Administrator Response: Correct the configuration file by removing the Configuration parameter.

GLD0188E Server is already listening on IP:

> ip_address and port: port_number. Ignoring listen: listen_arg.

Severity: Eventual Action

Explanation: The configuration file has more than one listen parameter specified to listen on the same IP address and port number.

System Action: The program continues.

Administrator Response: Correct the listen parameter in the configuration file by specifying a different IP address or port number. Then restart the server.

GLD0189I

Nonsecure communication is active for **IP:** *ip_address*, **nonsecure port**: port_number.

Severity: Information

Explanation: The LDAP server has activated the specified IP address ip address, and port for nonsecure communication. An *ip_address* of 'IN_ADDRANY' indicates that the server is listening on all active interface addresses.

System Action: The program continues.

Administrator Response: None.

GLD0190I

Secure communication is active for IP: ip_address, secure port: port_number.

Severity: Information

Explanation: The LDAP server has activated the specified IP address and port for secure communication. An ip_address of 'IN_ADDRANY' indicates that the server is listening on all active interface addresses.

System Action: The program continues.

Administrator Response: None.

GLD0191A

Nonsecure communication failed to activate for IP: ip address, port:

port_number; reason code: reason_code.

Severity: Immediate Action

Explanation: The LDAP server encountered an error when attempting to activate nonsecure communication. An ip_address of 'IN_ADDRANY' indicates that the server attempted to listen on all active interfaces. Preceding messages should indicate the cause of the

System Action: The program continues.

Operator Response: Correct the cause of the error.

GLD0192A

Secure communication failed to activate for IP: ip_address, port: port_number; reason code: reason_code.

Severity: Immediate Action

Explanation: The LDAP server encountered an error when attempting to activate secure communication. An ip_address of 'IN_ADDRANY' indicates that the server attempted to listen on all active interfaces. Preceding messages should indicate the cause of the error.

System Action: The program continues.

Operator Response: Correct the cause of the error.

GLD0193E

Debug level is not valid; reason code: reason_code, incorrect value 'bad token' detected at offset offset in string 'debug_string'.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the debug level on either the -d option on the command line invocation of the program or a console command is not valid. The reason code indicates why the value is not accepted. It can be:

- · -1 memory error in the server,
- -2 adjacent signs detected (for example, +-keyword),
- · -3 Debug value is too large,
- · -4 Debug value is not a valid hexadecimal,
- -5 Debug value is not a valid decimal or keyword,
- · -6 Off debug used with a sign (for example, +OFF, -OFF),
- -7 Debug value is only a sign (for example, +, -),
- -8 Debug value ends with a sign (for example, 8+, OFF-).

The token that caused the error is shown in the message along with its offset from the beginning of the debug string. The debug level is ignored.

System Action: The program continues.

Operator Response: If debug is needed, restart the

server with corrected command line options or specify a modify command from the operator's console to correctly specify the debug level.

Administrator Response: Correct the command line option and restart the server.

GLD0194I Successfully set debug level to

debug_level.

Severity: Information

Explanation: The debug level was set to the specified

value.

System Action: The program continues.

Administrator Response: None.

GLD0196A Attempt to accept() on IP: ip_address,

port port_number failed; errno errno

(errno_string).

Severity: Immediate Action

Explanation: The LDAP server received an error from system function **accept()** on the specified IP Address and port. Refer to the *z/OS C/C++ Run-Time Library Reference*, SC28-1663 for an explanation of the **errno** returned.

System Action: The LDAP server stops listening on the specified IP address and port.

Operator Response: Ensure TCP/IP is operating correctly. If the problem was caused by TCP/IP being stopped, then shutdown the LDAP Server and restart it once TCP/IP is restarted. If the problem was caused by another error, then correct the problem and shutdown the LDAP Server and restart it once the problem is corrected. If the problem can not be corrected, then save the diagnostic information and contact the system programmer.

Administrator Response: If the problem persists, contact the service representative.

GLD0197I The listen option in the configuration

file is being overridden on the command line: listen arg

Severity: Information

Explanation: The LDAP server has now been assigned to bind and listen based on the value supplied on the command line.

System Action: The program continues.

Administrator Response: None.

GLD0198l -p option on the command line is being ignored

Severity: Information

Explanation: The **-p** option is being ignored because there is at least one listen option specified in the configuration file or on the command line. The listen option(s) in the configuration file or on the command line will take precedence.

System Action: The program continues.

Administrator Response: None.

GLD0199I -s option on the command line is being ignored

Severity: Information

Explanation: The **-s** option is being ignored because there is at least one listen option specified in the configuration file or on the command line. The listen option(s) in the configuration file or on the command line will take precedence.

System Action: The program continues.

Administrator Response: None.

GLD0200E The value for the pcThreads option is

not numeric. The default

(default_pc_threads.) will be used.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the pcThreads option of the configuration file is not numeric. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

ille and restait the server.

GLD0201E The value for the pcThreads option is

out of range (min_threads, max_threads). The default (default_pc_threads) will be used.

Severity: Eventual Action

Explanation: The LDAP server encountered an error when processing the pcThreads parameter of the configuration file. The server will continue with the default value.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD02021 Program Call communication is active.

Severity: Information

Explanation: The LDAP server has activated support

for program call communication.

System Action: The program continues.

Administrator Response: None.

GLD0203A The Program Call initialization failed;

rc: return_code, reason code:

reason_code.

Severity: Immediate Action

Explanation: The LDAP server encountered an error during initialization of the program call environment. Communication with the LDAP server using program calls is not enabled.

System Action: PC initialization is terminated. Initialization continues for all other requested communication interfaces. If no communication interface is successfully initialized, the program ends.

Operator Response: Inform the System Programmer of the error.

Administrator Response: Restart the server after the system programmer has resolved the problem.

Programmer Response: Use the return code and reason code information to correct the problem. Possible return codes are:

- · 2 another server has already enabled PC communications.
- 3 cannot establish an ESTAEX exit. The reason code is the return code from the ESTAEX macro.
- 5 cannot obtain a system linkage index. The reason code is the return code from the LXRES macro.
- 6 cannot create a PC entry table. The reason code is the return code from the ETCRE macro.
- 7 cannot connect the PC entry table to the linkage table. The reason code is the return code from the ETCON macro.
- · 8 cannot create a named token. The reason code is the return code from the IEANTCR macro.
- 9 cannot make the address space non-swappable. The reason code is the return code from the SYSEVENT macro.
- 90 there is insufficent memory available.

GLD0204E Additional 'listen' configuration statements for Program Calls are ignored.

Severity: Eventual Action

Explanation: The LDAP server encountered multiple listen configuration statements for the program call environment. Only one is allowed. Additional program

call listen configuration statements are ignored.

System Action: The program continues.

Operator Response: Inform the Administrator of the

Administrator Response: Correct the configuration file to specify only one listen statement for the program call environment.

GLD0205A Program Call communication failed to activate; return code: return_code

Severity: Immediate Action

Explanation: The LDAP server encountered an error when attempting to activate program call communication. Preceding messages should indicate the cause of the error.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Examine other messages and correct any other reported errors. Restart the LDAP server.

GLD0206A Program Call (PC) support is not initialized because another server with PC support is already running.

Severity: Immediate Action

Explanation: There can be only one active LDAP server with program call (PC) support. Before initializing PC support, a new LDAP server checks that no active LDAP server on this system has already initialized PC support. If this is the case, PC support initialization in the new LDAP server is terminated.

System Action: PC initialization is terminated. Initialization continues for all other requested communication interfaces. If no communication interface is successfully initialized, the program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: If it is necessary to have PC support running in the new LDAP server, first remove the listen record for PC support from the configuration of the active LDAP server. Then restart both that LDAP server and the new LDAP server. The problem can be avoided by ensuring that only one LDAP server is configured to provide PC support.

GLD0207I backend_identifier manages the following suffixes:

Severity: Information

Explanation: The LDAP server is reporting the

suffixes managed by the backend.

System Action: The program continues.

Administrator Response: None.

GLD0208l Backend suffix: backend_suffix

Severity: Information

Explanation: The LDAP server is reporting the

suffixes managed by the backend.

System Action: The program continues.

Administrator Response: None.

GLD0209I End of suffixes managed by

backend_identifier.

Severity: Information

Explanation: The LDAP server has completed the list

of suffixes managed by the backend.

System Action: The program continues.

Administrator Response: None.

GLD0210I Modify command has been processed

successfully: modify text

Severity: Information

Explanation: A modify command has been processed

successfully.

System Action: The program continues.

Administrator Response: None.

GLD0211A The value specified for the adminDN option is 'cn=Anybody'. This is not

permitted.

Severity: Immediate Action

Explanation: The LDAP server encountered an error during processing of the adminDN option of the configuration file. The adminDN cannot be 'cn=Anybody'.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and restart the server.

GLD0212E

Cipher specification value is not valid; reason code: reason_code, incorrect value 'bad_token' detected at offset offset in string 'cipher_string'.

Severity: Eventual Action

Explanation: The LDAP server determined that the value specified for the cipher is not valid. The reason_code indicates why the value is not accepted. It can be:

- · -1 memory error in the server,
- -2 adjacent signs detected (for example, +-keyword),
- · -3 Cipher value is too large,
- · -4 Cipher value is not a valid hexadecimal,
- · -5 Cipher value is not a valid decimal or keyword,
- -6 Off Cipher used with a sign (for example, +OFF, -OFF),
- -7 Cipher value is only a sign (for example, +, -),
- -8 Cipher value ends with a sign (for example, 8+, OFF-).
- · -9 value specifies mask bits which are not allowed,
- · -10 an internal processing error was detected.

The token that caused the error is shown in the message along with its offset from the beginning of the cipher string. The cipher is ignored.

System Action: The program continues.

Operator Response: the cipher specification

corrected.

Administrator Response: If the reason code indicates an internal processing error, then contact IBM support and provide the message text. Otherwise, correct the cipher specification.

GLD0213I Cipher specifications mask set to cipher mask.

Severity: Information

Explanation: The LDAP server has assigned the cipher specification mask to the specified value based on the value read from the configuration file and the values supported by SSL.

System Action: The program continues.

Operator Response: None.

GLD0214E SSL cipher lowered from the specified value 'cipher_mask' to 'lowered'

Severity: Eventual Action

Explanation: The specified cipher value was not fully supported by SSL. A lower cipher value is used. This is normally caused by an incorrect specification of the cipher values or specification of a cipher value that is

not allowed in the region of the world where the LDAP server is running.

System Action: The program continues. If a cipher specification is needed, restart the server with the cipher specification corrected.

Administrator Response: None

GLD0215I End Successful Modify DN operation:

> **dn =>** *Distinguished:* newrdn =>new: deleteoldrdn => deleteOldRdn

Severity: Information

Explanation: The LDAP server successfully completed a Modify DN operation. This message displays the target DN, new RDN, and deleteOldRdn flag values, respectively, in the operation request. This information is provided to track the changes requested by Modify DN operations, because they may potentially affect many entries in the directory.

System Action: The program continues.

Operator Response: None. Administrator Response: None.

GLD0216I **End Successful Modify DN operation:**

dn => Distinguished; newSuperior => new

Severity: Information

Explanation: The LDAP server successfully completed a Modify DN operation. This message displays the target DN and new Superior DN respectively, in the operation request. This information is provided to track the changes requested by Modify DN operations, because they may potentially affect many entries in the

System Action: The program continues.

Operator Response: None. Administrator Response: None.

GLD0217I End Successful Modify DN operation:

> dn => Distinguished; control type => Control; control criticality => Control; control value (hexadecimal) => x'Control'

Severity: Information

Explanation: The LDAP server successfully completed a Modify DN operation. This message displays the target DN and the contents of any controls submitted in the operation request, including control type, control criticality, and control value (shown in hexadecimal format) for each control. This information is provided to track the changes requested by Modify DN operations, because they may potentially affect many entries in the directory.

System Action: The program continues.

Operator Response: None.

Administrator Response: None.

GLD0218E The value specified (value specified) for

the serverEtherAddr option is not correct. The default value of the CPU model and serial number will be used.

Severity: Eventual Action

Explanation: The LDAP server encountered an error during processing of the serverEtherAddr option in the configuration file. The default value will be used.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0219A Dynamic load of Kerberos DLL

dll_name failed. Errno: errno

(errno_string)

Severity: Immediate Action

Explanation: The server was unable to load the

Kerberos DLL.

System Action: The program ends.

Operator Response: None.

Administrator Response: Check your Kerberos

configuration.

GLD0220A Backend with name 'backend_name' already exists.

Severity: Immediate Action

Explanation: The LDAP server encountered an error processing the configuration file because the specified backend name is already in use by another backend. Note that 'cdbm1' is a reserved name and therefore cannot be used in the configuration file. Also, backend names are not case sensitive. For example, 'server1' and 'SERVER1' are duplicate values.

System Action: The program ends.

Administrator Response: Change the duplicate name to a unique string or remove the name from the database line in the configuration file, and then restart the server.

GLD0221A Cannot configure more than one sdbm backend.

Severity: Immediate Action

Explanation: The LDAP server encountered an error processing the configuration file beause a second sdbm database section was found. There can be at most one SDBM database section in the configuration file.

System Action: The program ends.

Administrator Response: Correct the configuration file so that it contains at most one sdbm database section and restart the server.

GLD0222E

Replica server on replica name on port port_number does not contain support for replicating Modify DN operations.

Severity: Eventual Action

Explanation: The replica server does not contain support for Modify DN operations with newSuperior and/or IBMModifyDNRealignDNAttributesControl and/or non-leaf Modify DN operations.

System Action: The program continues. Modify DN operations with newSuperior and/or IBMModifyDNRealignDNAttributesControl and/or non-leaf Modify DN operations will not be permitted.

Operator Response: None.

Administrator Response: Install the correct version of LDAP on the replica server or remove it from the replica collection.

GLD0223A Error %1\$d reported by Idap_api querying replica server version on

replica_name on port port_number.

Severity: Immediate Action

Explanation: An LDAP API call failed while querying the Replica server version.

System Action: The program continues. Modify DN operations with newSuperior and/or IBMModifyDNRealignDNAttributesControl and/or

non-leaf Modify DN operations will not be permitted.

Operator Response: Save the diagnostic information and contact the system programmer.

Administrator Response: If the problem persists, contact the service representative.

GLD0224A One or more replica servers do not implement the Modify DN function.

Severity: Immediate Action

Explanation: One or more replica servers do not implement, or can not be determined to implement, the requisite Modify DN function. However there is a committed Modify DN operation on the master server which can not be replicated to one or more of the replicas. This state can result in continuing replication failures (until the situation is rectified) which will result in diverging directory contents between master server and replica server(s).

System Action: The program continues. The replication operation will not be performed.

Operator Response: None.

Administrator Response: Make sure that all replicas have the required level of the LDAP server.

GLD02251 Ignoring suffix: backend_suffix defined for exop backend.

Severity: Information

Explanation: The LDAP server encountered a suffix for a backend of type exop in the configuration file. An exop backend

System Action: The program continues.

Administrator Response: None required, correct the

configuration file.

GLD0226E The 'attribute' parameter is no longer supported. All 'attribute' lines will be

ignored.

Severity: Eventual Action

Explanation: The attribute parameter is now obsolete and is being ignored. Attribute definitions should be provided as documented for each server backend type. See the LDAP Server Administration and Use for more information.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0227E The 'object class' parameter is no longer supported. All 'objectclass'

lines will be ignored.

Severity: Eventual Action

Explanation: The objectclass parameter is now obsolete and is being ignored. Objectclass definitions should be provided as documented for each server backend type. See the LDAP Server Administration and Use for more information.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0228E The 'verifySchema' parameter is no longer supported. This line will be

ignored.

Severity: Eventual Action

Explanation: The verifySchema parameter is now obsolete and is being ignored. The verifySchema parameter was only applicable for attribute and objectclass parameters, which are no longer supported. Note that schema loaded into the LDAP server for TDBM is always verified, regardless of the presence of the verifySchema parameter.

System Action: The program continues.

Administrator Response: Correct the configuration file and restart the server.

GLD0229E The activity logfile could not be

Severity: Eventual Action

Explanation: The LDAP server could not open the activity logfile. This is due to an error in specifying the logfile configuration parameter.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

GLD0230A A self-signed SSL certificate cannot be validated because it is not in the SSL key database or key ring.

Severity: Immediate Action

Explanation: A self-signed certificate cannot be validated because it is not in the key database file or key ring.

System Action: The program continues. The request fails.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Add the self-signed certificate to the key database file or key ring.

GLD0231A

The SSL Peer application sent a certificate for which the certification authority is not in the SSL key database or key ring.

Severity: Immediate Action

Explanation: The key database file or key ring does not contain a certificate for the certification authority.

System Action: The program continues. The request fails.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Obtain the certificate for the certification authority and add it to the key database file or key ring.

GLD0232A

The SSL server certificate is not compatible with the negotiated cipher suite.

Severity: Immediate Action

Explanation: The certificate key is not compatible with the negotiated cipher suite. The server certificate must have an RSA key while the client certificate may have an RSA or DSA key. This error can also occur if the

client certificate has a DSA key but the server does not support DSA keys, the server key usage certificate extension does not allow key encipherment, or the client key usage certificate extension does not allow digital signature. For the 40-bit export ciphers, the server key usage certificate extension must allow digital signature.

System Action: The program continues. The request fails.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Specify a certificate with the appropriate key type and key usage.

GLD0233A SSL detected an error while validating a certificate.

Severity: Immediate Action

Explanation: An error is detected while validating a certificate. This error can occur if a root CA certificate is not found in the key database file or key ring or if the certificate is not marked as a trusted certificate.

System Action: The program continues. The request fails.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Verify that the root CA certificate is in the key database or key ring and is marked as trusted. Check all certificates in the certification chain and verify that they are trusted and are not expired.

GLD0234A A memory allocation error occurred in SSL processing.

Severity: Immediate Action

Explanation: The LDAP server is unable to allocate the necessary storage for SSL processing.

System Action: Communication interfaces that require SSL will not be activated. Communication interfaces that support both secure and non-secure communication will be activated for non-secure communication only. If no interfaces are activated then the program will end.

Operator Response: Increase the storage for the LDAP server and restart the server.

GLD0235A

The maximum connections limit could not be set to

'new_max_connection_limit'. The setrlimit() call failed with errno='errno' (errno_string). The limit remains unchanged at

'existing_max_connection_limit'.

Severity: Immediate Action

Explanation: The LDAP server was unable to change

the number of connections allowed to the limit specified by the maxconnections parameter in the configuration file. The limit is affected by the number of file descriptors currently open for the process and the system limit of 65,535.

System Action: The program continues with the existing maximum connections value.

Administrator Response: Determine the reason for the error. The errno value set by setrlimit() is provided to assist in identifying the error. Correct the maxconnections value in the configuration file and start the server again.

GLD0236I

IP: ip_address, port port_number has been stopped.

Severity: Information

Explanation: The LDAP server has stopped communication on the specified IP address and port. No new connections will be accepted on the specified interface.

System Action: The LDAP server stops communicating on the specified IP address and port.

Operator Response: If communication is desired on the specified interface, then stop the LDAP server, if it is still running. Restart the LDAP server once it is stopped.

Administrator Response: None.

GLD0237I

Program Call Interface has been stopped.

Severity: Information

Explanation: The LDAP server has stopped communication on the program call interface. No new connections will be accepted on the specified interface.

System Action: The LDAP server stops communicating on the specified interface.

Operator Response: If communication is desired on the specified interface, then stop the LDAP server, if it is still running. Restart the LDAP server once it is stopped.

Administrator Response: None.

GLD0238A

Backend type (type) and DLL do not match.

Severity: Immediate Action

Explanation: The LDAP server encountered an error processing the configuration file. The specified database type and the DLL to load do not match. One possible

reason is a typographical error in the configuration file.

System Action: The program continues through configuration processing but the configuration will fail and the server will not start.

Administrator Response: Check the database lines in the configuration file for the specified backend type and verify the correct DLL is being loaded. Correct the configuration file and restart the server. If the problem persists, contact the service representative.

GLD0239A

ibmdirectoryversion attribute not found in Root DSE of replica server on replica_name on port port_number.

Severity: Immediate Action

Explanation: It is not possible to determine if the replica server contains support for Modify DN operations with newSuperior and/or

IBMModifyDNRealignDNAttributesControl and/or non-leaf Modify DN operations or support for ibm-entryuuid attributes.

System Action: The program continues. ibm-entryuuid attributes will not be replicated. Modify DN operations with newSuperior and/or

IBMModifyDNRealignDNAttributesControl and/or non-leaf Modify DN operations will not be permitted.

Administrator Response: Install a conformant version of LDAP on the replica server or remove it from the replica collection.

GLD0240A

ibmdirectoryversion attribute has no value in Root DSE of replica server on replica_name on port port_number.

Severity: Immediate Action

Explanation: It is not possible to determine if the replica server contains support for Modify DN operations with newSuperior and/or

IBMModifyDNRealignDNAttributesControl and/or non-leaf Modify DN operations or support for ibm-entryuuid attributes.

System Action: The program continues. ibm-entryuuid attributes will not be replicated. Modify DN operations with newSuperior and/or

IBMModifyDNRealignDNAttributesControl and/or non-leaf Modify DN operations will not be permitted.

Administrator Response: Install a conformant version of LDAP on the replica server or remove it from the replica collection.

Administration messages (2000)

GLD2086E Encrypt all passwords that are presently in CLEAR format (yes/no)?

Severity: Eventual Action

Explanation: Prompt a user with this message to give her/him a chance to change their mind If the response was 'yes' (or 'y' or 'Y'), the program will continue, otherwise stopped.

System Action: The program continues or ends

based upon input.

Operator Response: None. Administrator Response: None

GLD2091A No base is defined.

Severity: Immediate Action

Explanation: A base must be defined either on the command line or through the LDAP_BASEDN

environment variable.

System Action: The program ends.

Operator Response: None.

Administrator Response: Define a base and try the

program again.

GLD2092A db2pwden ends without encrypting passwords.

Severity: Immediate Action

Explanation: The user changed their mind or an error occurred which caused the db2pwden program to end without encrypting passwords.

System Action: The program ends.

Operator Response: None.

Administrator Response: If the user wants to try the program again, correct any errors identified and try

again.

GLD2100A Memory allocation failed.

Severity: Immediate Action

Explanation: The progam has used all of the available

storage.

System Action: The program ends.

Operator Response: None.

Administrator Response: Increase the region size

and try again.

GLD2101A Only supported mechanisms are EXTERNAL, GSSAPI, CRAM-MD5, and **DIGEST-MD5**

Severity: Immediate Action

Explanation: The progam does not support the

requested authentication mechanism.

System Action: The program ends.

Operator Response: None.

Administrator Response: Specify EXTERNAL,

GSSAPI, CRAM-MD5, or DIGEST-MD5.

GLD2102A Scope should be base, one, or sub

Severity: Immediate Action

Explanation: The progam does not support the

requested scope mechanism.

System Action: The program ends.

Operator Response: None.

Administrator Response: Specify either base, one or

sub.

Error error_code and reason GLD2103A

reason_code reported by SSL Client

initialization.

Severity: Immediate Action

Explanation: The LDAP utility was unable to complete

initialization required for secure communications.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2104A Error reported by LDAP client initialization.

Severity: Immediate Action

Explanation: The LDAP utility was unable to complete

initialization required for communications.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2105A Error error_code reported binding to LDAP server.

Severity: Immediate Action

Explanation: The LDAP utility was unable to bind to

the server.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2106A Error error_code reported modifing DN

distinguished_name.

Severity: Immediate Action

Explanation: The LDAP utility was unable to add the

ibm-entryuuid attribute to the LDAP entry. **System Action:** The program terminates. **Operator Response:** Contact the service

representative.

GLD2107A Error error_code reported parsing Idap

results.

Severity: Immediate Action

Explanation: The LDAP utility was unable to add the

ibm-entryuuid to the LDAP entry.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2108A Error error_code reported by

Idap_search.

Severity: Immediate Action

Explanation: The LDAP search failed.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2109A Error error_code reported by

ldap_get_dn.

Severity: Immediate Action

Explanation: The LDAP utility was unable to get the

DN.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2110A Error error_code reported by

Idap_first_entry.

Severity: Immediate Action

Explanation: An LDAP internal error occurred.

System Action: The program terminates.

Operator Response: Contact the service

representative.

GLD2111I Idapadduuids added ibm-entryuuids to

matches entries.

Severity: Information

Explanation: The **Idapadduuids** utility added ibm-entryuuids to the specified number of entries.

System Action: None.

Operator Response: None.

GLD2112I command_name processed matches

entries.

Severity: Information

Explanation: The utility processed the specified

number of entries.

System Action: None.

Operator Response: None.

GLD2113A A bind DN is required.

Severity: Immediate Action

Explanation: The LDAP utility will not function unless

a bind DN is specified.

System Action: The program terminates.

Operator Response: Specify a bind DN.

GLD2114A A basedn must be specified.

Severity: Immediate Action

Explanation: The LDAP utility will not function unless

a base DN is specified.

System Action: The program terminates.

Operator Response: Specify a base DN either with the **-b** option or with the **LDAP_BASEDN** environment

variable.

GLD2116A A username is required when doing a

DIGEST-MD5 bind.

Severity: Immediate Action

Explanation: The LDAP utility will not function unless

a user name is specified.

System Action: The program terminates.

Operator Response: Specify a user name with the -U

option.

GLD2117A Debug value is not valid.

Severity: Immediate Action

Explanation: The LDAP utility will not function because the specified debug value is not valid.

System Action: The program terminates.

Operator Response: Specify a valid debug value with

the -d option.

GLD2118A Incorrect DN syntax.

Severity: Immediate Action

Explanation: The LDAP utility was unable to bind to

the server.

System Action: The program terminates.

Operator Response: Correct the DN syntax.

GLD2119A Incorrect LDAP server name or LDAP

server is not available.

Severity: Immediate Action

Explanation: The LDAP utility was unable to bind to

the server.

System Action: The program terminates.

Operator Response: Specify the correct LDAP server

name or start the LDAP server.

GLD2120A Credentials are not valid for the

specified LDAP server.

Severity: Immediate Action

Explanation: The LDAP utility was unable to bind to the LDAP server because the specified credentials were

not accepted by the LDAP server.

System Action: The program terminates.

Operator Response: Specify the correct credentials

for the LDAP server name.

GLD2121A Base DN not found on LDAP server.

Severity: Immediate Action

Explanation: The LDAP utility was unable to perform the operation because the specified base DN was not

found by the LDAP server.

System Action: The program terminates.

Operator Response: Specify a base DN that exists on

the LDAP server.

GLD2122A Bind to LDAP server failed, DN not

found and default referral not

permitted.

Severity: Immediate Action

Explanation: The LDAP utility was unable to bind to

the LDAP server because the specified bind DN was not found by the LDAP server and the use of default referral is not permitted.

System Action: The program terminates.

Operator Response: Specify a DN that exists on the

LDAP server.

GLD2123A Object not found on LDAP server.

Severity: Immediate Action

Explanation: The LDAP server was unable to find an object that matched the specification and that did not

have an ibm-entryuuid attribute.

System Action: The program terminates.

Operator Response: Use Idap_search to make sure the object you are attempting to modify exists on the LDAP server. Specify ibm-entryuuid as the attribute

Idap_search should return.

GLD2124A LDAP server unwilling to perform specified operation.

Severity: Immediate Action

Explanation: The LDAP server is not permitted to

perform the operation.

System Action: The program terminates.

Operator Response: Check the options passed to the

utility.

GLD2125A Error in search filter.

Severity: Immediate Action

Explanation: The search filter is not formed with the correct syntax or there is an incorrect character in the

specification.

System Action: The program terminates.

Operator Response: Check the search filter.

GLD3001I time command_name: number entries have been processed.

Severity: Information

Explanation: The program has processed the specified number of entries. Note that the program might have encountered errors during this processing. If so, processing of some entries may not have completed successfully. Additional messages are issued to indicate these errors.

System Action: The program continues.

Administrator Response: None.

GLD3002I command_name has completed

successfully.

Severity: Information

Explanation: The command has ended without

encountering any errors.

System Action: The program ends.

Administrator Response: None.

GLD3003I command_name has failed.

Severity: Information

Explanation: The command has ended after

encountering an error.

System Action: The program ends.

Administrator Response: Use the information in any error messages issued by the command to fix the problem. Then try the command again.

problem them by the command again.

GLD3004l The check step of command_name has completed successfully.

Severity: Information

Explanation: The check step of the **Idif2tdbm** program has ended without encountering any errors.

System Action: The program continues to process other requested steps.

Administrator Response: None.

GLD3005I The check step of command_name has failed.

Severity: Information

Explanation: The check step of the **Idif2tdbm** program has ended after encountering errors.

System Action: The program ends.

Administrator Response: Use the information in any error messages issued by the command to fix the problem. Then try the command again.

GLD3006l The prepare step of command_name has completed successfully.

Severity: Information

Explanation: The prepare step of the **Idif2tdbm** program has ended without encountering any errors.

System Action: The program continues to process

other requested steps.

Administrator Response: None.

GLD3007I The prepare step of command_name has failed.

Severity: Information

Explanation: The prepare step of the **Idif2tdbm** program has ended after encountering errors. Existing data in the output files from a previous invocation of the prepare step have been deleted.

System Action: The program continues to process the check step if it has been requested. Otherwise, the program ends.

Administrator Response: Use the information in any error messages issued by the command to fix the problem. Then try the command again.

GLD3008l The schema modify part of the load step of command_name has completed successfully.

Severity: Information

Explanation: If one or more schema input files are specified with the **-s** or **-v** options, the first part of the load step of the **Idif2tdbm** program is to modify the schema in the database, using the schema input files. The schema in the database has been successfully updated.

System Action: The program continues to process the load step.

Administrator Response: None.

GLD3009I The load step of command_name has successfully submitted the DB2 load utility jobs.

Severity: Information

Explanation: The load step of the **Idif2tdbm** program submits the DB2 Load utility jobs created during the prepare step to load the new entries into the database. The load jobs have been successfully submitted. **Note:** This message does not indicate that the load jobs have terminated successfully. The processing of the load jobs by DB2 is outside the scope of the **Idif2tdbm** program. You must review the output generated by each load job to determine if it was successful.

System Action: The program ends.

Administrator Response: Review the output of each load to determine if it was successful. If not, use the information in the description of the Idif2tdbm command in z/OS Security Server LDAP Server Administration and Use Guide to determine how to proceed. Note: Do not run the Idif2tdbm program again because this can add duplicate data to the database.

GLD3010I The load step of command_name has failed.

Severity: Information

Explanation: The load step of the **Idif2tdbm** program has ended after encountering errors. However, if one or more schema input files were specified with the -s or -v options, the first part of the load step, which modifies the schema in the database, may have succeeded. If so, a message indicating this will precede this message.

System Action: The program ends.

Administrator Response: Use the information in any error messages issued by the command to fix the problem and determine how to proceed. If an error message was issued indicating a failure while submitting JCL, do not run the Idif2tdbm program again because this can add duplicate data to the database. Instead, use the information in the description of the Idif2tdbm command in z/OS Security Server LDAP Server Administration and Use Guide to determine how to proceed. If the JCL message was not issued, try the Idif2tdbm command again. Do not specify the -s or -v options to modify the schema if that part of the load step was successful.

GLD3011I The command_name status file file name cannot be written.

Severity: Information

Explanation: The Idif2tdbm was not able to write the new status file after completing processing of the requested steps. The current status file may have been deleted. All other output files, such as the load and JCL datasets produced by the prepare step, have not been affected. This will likely result in warning messages the next time that Idif2tdbm is invoked for this output file.

System Action: The program ends.

Administrator Response: Use the information in the error messages issued by the command to fix the problem. If the error messages indicate that some processing step failed, run the command again. If Idif2tdbm issues warning messages and a prompt, be sure that all the conditions in the warning messages are acceptable before responding to continue processing.

GLD3012I

command name has terminated because there are no entries to process.

Severity: Information

Explanation: There are no entries for the program to process. For the Idif2tdbm program, there are no entries in the LDIF input files. For the tdbm2ldif program, there are no entries in the database.

System Action: The program ends.

Administrator Response: For Idif2tdbm, ensure that the LDIF input files contain entries to add to the database. For tdbm2ldif, check that the database contains entries to unload. Then try the program again.

GLD3013A

program_name: line line_number: incorrect configuration line: option takes parameters values.

Severity: Immediate Action

Explanation: The TDBM backend cannot be configured because of the specified error in the configuration file.

System Action: The program continues. If all required parameters are not set correctly, the configuration will

Administrator Response: Correct the configuration file and try again.

GLD3015A

program_name: line line_number: incorrect configuration line: unrecognized keyword.

Severity: Immediate Action

Explanation: The TDBM backend cannot be configured because the specified line in the configuration file contains an unrecognized keyword.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and try again.

GLD3016A

keyword parameter is missing from configuration file.

Severity: Immediate Action

Explanation: The TDBM backend cannot be configured because the configuration file is missing a required parameter.

System Action: The program continues. If all required parameters are not set correctly, the configuration will fail and the server will not start. If the configuration parameter that failed is associated with a specific database backend, the server may start, but the backend may not be available. If a configuration parameter is ignored, the server may not operate as expected.

Administrator Response: Correct the configuration file and try again.

GLD3017A Unable to connect to the database, rc = return_code.

Severity: Immediate Action

Explanation: The LDAP server or LDAP utility cannot connect to DB2.

Connect to DB2.

System Action: The program continues. The TDBM

backend ends.

Operator Response: Ensure that DB2 is started and

running properly.

GLD3018A The dsnaoini CLI initialization file was not specified.

Severity: Immediate Action

Explanation: The LDAP server or LDAP utility cannot complete initialization of the DB2 interface because dsnaoini was not specified in the slapd.conf file, there is no DD specified for DSNAOINI in the JCL, or the DSNAOINI environment variable is not set.

System Action: The program continues. The TDBM backend ends.

Administrator Response: Specify dsnaoini in the slapd.conf file, the DD DSNAOINI in the JCL, or set the DSNAOINI environment variable.

GLD3019E Entry 'distinguished_name' already exists.

Severity: Eventual Action

Explanation: The request cannot be completed because an attempt is being made to add an entry that already exists in the database.

System Action: The program continues. The request fails.

Administrator Response: Correct the request and try again.

GLD3020E Entry 'distinguished_name' violates the schema definition.

Severity: Eventual Action

Explanation: The request cannot be completed because an attempt is being made to add an entry that

does not match the schema definition.

System Action: The program continues. The request fails.

Administrator Response: Correct the request and try again.

GLD3021E Parent entry does not exist for entry 'distinguished name'.

Severity: Eventual Action

Explanation: The request cannot be completed because no parent entry exists for the entry being added

System Action: The program continues. The request

fails.

Administrator Response: Correct the request and try

again.

GLD3033I The LDAP server will operate in multi-server mode.

Severity: Information

Explanation: The LDAP server or LDAP utility will operate in multi-server mode. Multiple concurrent servers may be running on a given z/OS image, and multiple concurrent servers may be running on multiple z/OS images in a Parallel Sysplex, all of which use the same LDAP DB2 database. Note that no replication may be performed by any of these servers. If replication is desired, only one server instance may be running which uses a given LDAP DB2 database, and the server must be configured to run in single-server mode.

System Action: The program continues.

Administrator Response: None.

GLD3036E Multiserver keyword argument is 'n' or 'N', but the LDAP server will operate in multi-server mode.

Severity: Eventual Action

Explanation: The multiserver keyword argument found in the server configuration file was either 'n' or 'N'. However, because both sysplexServerName and sysplexGroupName keywords have valid arguments, the LDAP server or LDAP utility must operate in multi-server mode.

System Action: The program continues.

Administrator Response: Correct the configuration

file and restart the server.

The LDAP program found no TDBM GLD3039A database.

Severity: Immediate Action

Explanation: The configuration file does not contain any TDBM database specifications. The program cannot continue because it cannot find the appropriate database with which to work.

System Action: The program ends.

Administrator Response: Ensure that the database needed by the program is specified in the configuration file. Then try the program again.

GLD3040A Overlapping TDBM backend suffixes found in configuration file: suffixes

'first_overlapping_suffix' and 'second_overlapping_suffix' overlap.

Severity: Immediate Action

Explanation: The presence of overlapping suffixes was detected in the server configuration file for a TDBM backend. Overlapping suffixes (suffixes for which the hierarchy is identical to the extent of the shorter of the two) may not be specified. Eliminate the overlap in suffixes and restart the server.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Eliminate the overlap in TDBM suffixes and restart the server.

GLD3041E Parent of new entry is a referral object. Cannot add new entry

'distinguished_name'.

Severity: Eventual Action

Explanation: The request cannot be completed because an entry cannot be added directly below a referral object. The entry must be added to the namespace below the actual entry which is referenced by the referral object.

System Action: The program continues. The request fails.

Administrator Response: Direct the request to the correct location in the namespace and try again.

GLD3042A The LDAP program encountered an error during configuration.

Severity: Immediate Action

Explanation: An LDAP program encountered an error while processing the configuration file. The program cannot continue. See additional messages for more information about the error encountered.

System Action: The program ends.

Administrator Response: Examine the additional messages. Correct the configuration file and try the program again.

GLD3043A The LDAP program found incomplete database information.

Severity: Immediate Action

Explanation: An LDAP program encountered an empty list of TDBM-specific information. The program cannot continue.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Correct the configuration file and restart the server.

GLD3045A The LDAP program does not support this encryption method.

Severity: Immediate Action

Explanation: The LDAP configuration file contains the keyword pwEncryption with an incorrect value. Correct values are none, crypt, MD5, SHA, or DES.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Correct the configuration file and restart the server.

GLD3046A OCSF setup failed, encryption method 'method' is not available.

Severity: Immediate Action

Explanation: The pwEncryption value specified in the configuration file requires the OCSF product to be installed and available on your system.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Verify OCSF setup and try again or change the pwEncryption value in the configuration file to a method that does not need OCSF (none or crypt).

GLD3047E OCSF setup failed, but pwEncryption method 'method' is available.

Severity: Eventual Action

Explanation: The pwEncryption value specified in the configuration file does not require the OCSF product to

be installed and available on your system, but any value already encrypted in MD5, SHA, or DES will not compare correctly on a bind.

System Action: The program continues.

Administrator Response: If OCSF is needed, verify OCSF setup and then stop and start the LDAP program again.

GLD3048A DES key label not available, encryption method 'method' is not available.

Severity: Immediate Action

Explanation: The pwEncryption value specified in the configuration file requires the OCSF product and the ICSF product to be installed and available on your system. It also requires a valid CKDS and the key corresponding to the DES key label in the configuration file to be available on your system.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Verify OCSF, ICSF, and CKDS setup and try again or change the pwEncryption value in the configuration file to a method that does not need OCSF, ICSF, or CKDS.

GLD3049A The DES key label specified with pwEncryption in the configuration file is too long.

Severity: Immediate Action

Explanation: The DES key label specified with pwEncryption in the configuration file can be a maximum of 64 characters long.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Set up a key with a key label with a valid length and try again.

GLD3050A The format of pwEncryption with DES is incorrect in the configuration file.

Severity: Immediate Action

Explanation: The format of pwEncryption with DES should be **DES**: *keylabel*.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Format the pwEncryption value correctly and restart the server.

GLD3053E

Attention: Idif2tdbm has detected a request for the requested_step processing step but the previous step may not have completed successfully.

Severity: Eventual Action

Explanation: This invocation of Idif2tdbm uses the same output datasets (specified using the -o option) as a previous invocation, hence is seen as a continuation of processing of the earlier invocation. The Idif2tdbm program compares the current status in the status file with the processing steps specified on the command to determine if any processing step may have been skipped. Normally, the check (-c), prepare (-p), and load (-I) steps are performed in this order. If the program cannot determine that this order is being followed, this warning is issued. For example, if the current status is L (load) and only -p (prepare) is specified on the command, this warning message is issued to indicate that the step before the **-p** step (that is, the check step) may have been skipped. The status for the earlier invocation of Idif2tdbm is stored in the STATUS record in the status file, hlq.BULKLOAD.JCL(STATUS), where hlq is the value specified in the **-o** option in the command.

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is **1** (yes), processing continues. If the response is **0** (no), or any character except **1**, processing ends.

Administrator Response: Use the information in the message, the current status in the STATUS record in the status file, and the processing steps specified on the command to determine if any necessary steps have been skipped. If not, respond 1 (yes) to the prompt to allow processing to continue. Otherwise respond 0 (no), or any character except 1, to stop processing and then try the program again, specifying the appropriate processing steps on the command. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3054E

Attention: the specified list of schema files does not match the list of schema files specified in a previous invocation of Idif2tdbm.

Severity: Eventual Action

Explanation: This invocation of **Idif2tdbm** uses the same output datasets (specified using the **-o** option) as a previous invocation, hence is seen as a continuation of processing of the earlier invocation. The **Idif2tdbm** program has detected that the list of schema files specified in the **-s** and the **-v** options of this invocation is different than on the earlier invocation in one of two ways: the lists contain different names or the order of

names in the lists is different. Since LDIF entries that were checked or prepared using the earlier list of schema file names may no longer be valid if processed using the new list of the schema file names, this warning message is issued. The schema file names for the the earlier invocation of Idif2tdbm are stored in the SCHEMA records in the status file, hla.BULKLOAD.JCL(STATUS), where hla is the value specified in the **-o** option in the command.

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is 1 (yes), processing continues. If the response is 0 (no), or any character except 1, processing ends.

Administrator Response: Determine whether processing the LDIF entries using the new list of schema file names is acceptable. If so, respond 1 (yes) to the prompt to allow processing to continue. Otherwise respond 0 (no), or any character except 1, to stop processing. Then try the program again, either specifying the earlier value for the -s and -v options, or specifying the new value for the options but changing the requested processing steps to redo the prior steps. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3055E

Attention: the specified list of LDIF files does not match the list of LDIF files specified in a previous invocation of Idif2tdbm.

Severity: Eventual Action

Explanation: This invocation of Idif2tdbm uses the same output datasets (specified using the -o option) as a previous invocation, hence is seen as a continuation of processing of the earlier invocation. The Idif2tdbm program has detected that the list of LDIF input file names specified in the -i and the -e options of this invocation is different than on the earlier invocation in one of two ways: the lists contain different names or the order of names in the lists is different. Since LDIF entries that were checked or prepared using the earlier list of LDIF input files may no longer be valid for the new list of LDIF input files, this warning message is issued. The LDIF input file names for the earlier invocation of Idif2tdbm are stored in the LDIFFILE records in the status file, hlg.BULKLOAD.JCL(STATUS), where *hlq* is the value specified in the **-o** option in the command.

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is 1 (yes), processing continues. If the response is 0 (no), or any character except 1, processing ends.

Administrator Response: Determine whether processing should continue using the new list of LDIF

input file names. If so, respond 1 (yes) to the prompt to allow processing to continue. Otherwise respond 0 (no), or any character except 1, to stop processing. Then try the program again, either specifying the earlier values for the -i and -e options, or specifying the new values for the options but changing the requested processing steps to redo the prior steps. Note that these options are processed in the order they are specified in the command. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3056E

Attention: Idif2tdbm will overwrite all files that exist in the output datasets during the prepare step.

Severity: Eventual Action

Explanation: During the prepare step, the ldif2tdbm program writes files to the output datasets. Since the Idif2tdbm program will overwrite the updates from an earlier Idif2tdbm program invocation, this warning message is issued. To determine whether to issue this message, the Idif2tdbm program reads the STATUS record in the status file, hlq.BULKLOAD.JCL(STATUS), where *hlq* is the value specified using the **-o** option. If the STATUS record has a value of P or L, the Idif2tdbm program recognizes that an earlier invocation has written to the output datasets and issues this message.

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is 1 (yes), processing continues. If the response is **0** (no), or any character except 1, processing ends.

Administrator Response: If the output from the previous invocation can be discarded, respond 1 (yes) to the prompt to allow processing to continue. Otherwise respond 0 (no), or any character except 1, to stop processing and then try the program again, specifying a new valid dataset high level qualifier using the -o option. Make sure that all required datasets with the new dataset high level qualifier are allocated before running the program again. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3057E

Attention: attempting to re-run the load step, after Idif2tdbm has successfully completed a load step.

Severity: Eventual Action

Explanation: During the load step (-I), the Idif2tdbm program submits DB2 load utility jobs and modifies the schema if it is specified (using the -s or -v options). Once the jobs are submitted, the Idif2tdbm program does not have any control over the DB2 load utility jobs. Moreover, resubmitting the DB2 load utility jobs with the same data will result in duplicate data in the database. Hence, the **Idif2tdbm** program should not attempt the

load step again, without preparing new LDIF input files. The **Idif2tdbm** program issues this message when the STATUS record in the status file.

hlq.BULKLOAD.JCL(STATUS), where hlq is the value specified in the -o option in the command, has a value of L.

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is 1 (yes), processing continues. If the response is 0 (no), or any character except 1, processing ends.

Administrator Response: If any of the DB2 load utility jobs submitted by an earlier invocation of the Idif2tdbm program completed successfully, respond 0 (no), or any character except 1, to stop processing. Review the output of each load to determine whether it was successful, refer to the DB2 Utility Guide and Reference to correct any problems, and then manually resubmit the DB2 load utility jobs that failed. If all the data prepared by an earlier invocation must be reloaded, respond 1 (yes) to the prompt to allow processing to continue. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3059E Continue despite previous attention messages? (0=no/1=yes)

Severity: Eventual Action

Explanation: The **Idif2tdbm** program is requesting permission from the administrator to continue despite the warning messages that were previously issued.

System Action: If the administrator responds 1, the program continues; otherwise, the program ends.

Administrator Response: The administrator must respond to the prompt, answering either 1 (yes) to allow the program to continue or 0 (no), or any character except 1, to stop processing.

GLD3062l Password encryption method 'method' is enabled in the TDBM backend.

Severity: Information

Explanation: The LDAP server will encrypt passwords in the TDBM backend using the encryption method specified.

System Action: The program continues.

Administrator Response: None.

GLD3063A An unrecognized status, *value*, encountered in status file.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program has detected an unrecognized value for the STATUS record in the status

file. The status file is hlq.BULKLOAD.JCL(STATUS), where hlq is the value specified in the **-o** option in the command.

System Action: The program ends.

Administrator Response: Correct the value in the STATUS record. Acceptable values are **N** for none, **C** for check, **P** for prepare, or **L** for Load, depending on the last step that **Idif2tdbm** successfully completed.

GLD3064A No value is specified for option: option.

Severity: Immediate Action

Explanation: The program was invoked with an option that is missing a value. The option displayed in the message must have a value.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again with the proper options and values.

GLD3065A An unrecognized parameter is specified: parameter.

Severity: Immediate Action

Explanation: The program was invoked with a parameter that it does not support. The parameter is displayed in the message.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again with the proper parameters and values.

GLD3066A Unrecognized value, value, for parameter, parameter.

Severity: Immediate Action

Explanation: The program was invoked with a value that is not valid for a parameter or option. The parameter and value are displayed in the message.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again with the proper parameters and values.

GLD3067A Parameter, parameter1, requires parameter, parameter2.

Severity: Immediate Action

Explanation: The program was invoked with a parameter that cannot be specified without also specifying a second parameter. The second parameter is missing. The specified parameter and the missing parameter that it requires are displayed in the message.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again, making sure that all proper parameter combinations are specified.

GLD3068A A required parameter was not specified.

Severity: Immediate Action

Explanation: The program has been invoked without a

required parameter.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again, making sure that all required parameters are specified.

GLD3069A A required parameter, parameter, was not specified.

Severity: Immediate Action

Explanation: The program has been invoked without a required parameter. The missing parameter is displayed in the message.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again, making sure that all required parameters are specified.

GLD3071A Unrecognized keyword, keyword, in status or system file.

Severity: Immediate Action

Explanation: The Idif2tdbm program detected a keyword that is not valid in the status file or the system file. The status file is hlq.BULKLOAD.JCL(STATUS) and the system file is hlq.BULKLOAD.JCL(SYSTEM), where hlq is the value specified in the -o option on the ldif2tdbm command.

System Action: The program ends.

Administrator Response: Use the information in the record to locate the keyword in the status or system file. Either correct the keyword or remove the record. Then try the program again.

GLD3072A An error occurred while reading line line_number of file file_name.

Severity: Immediate Action

Explanation: The program encountered an error while trying to read the line specified in the message.

System Action: The program ends.

Administrator Response: Ensure that the file exists and that the specified line can be read. Then try the program again.

GLD3073A

A second DN record is found in an LDIF entry at line line number of LDIF file file name.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program encountered an error while processing an entry in an LDIF file. The entry has more than one DN record. A DN record is a record that begins with dn:.

System Action: The program continues check processing if this entry is not the first entry in the first LDIF file and if the check step has been requested; otherwise, the program ends.

Administrator Response: Use the information in the message to locate the entry in the LDIF file and remove one of the DN records. Then try the program again.

GLD3074A

A DN record without a value is found in an LDIF entry at line line_number of LDIF file file_name.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program encountered an error while processing an entry in an LDIF file. The entry has a DN record without any value. A DN value is required for each entry. A DN record is a record that begins with dn:.

System Action: The program continues check processing if this entry is not the first entry in the first LDIF file and if the check step has been requested; otherwise, the program ends.

Administrator Response: Use the information in the message to locate the entry in the LDIF file and add a value to the DN record. Then try the program again.

GLD3075A

A non-comment record outside an entry is found at line line_number of LDIF file file name.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program encountered an error while processing an LDIF file. The specified record is not part of an entry.

System Action: The program continues check processing if this entry is not the first entry in the first LDIF file and if the check step has been requested; otherwise, the program ends.

Administrator Response: Use the information in the message to locate the record in the LDIF file. Either remove the record or move it so that it is within an entry. Then try the program again.

GLD3076A A value that cannot be processed is found in an LDIF entry at line

line number of LDIF file file name.

Severity: Immediate Action

Explanation: The Idif2tdbm program encountered an error while processing an entry in an LDIF file. The value that cannot be processed could be the type or value part of the record.

System Action: The program continues check processing if this entry is not the first entry in the first LDIF file and if the check step has been requested; otherwise, the program ends.

Administrator Response: Use the information in the message to locate the record in the LDIF file. Ensure that both the type and value are valid. Then try the program again.

GLD3077A

A required change indication line is missing in an LDIF modify entry at line line number of LDIF file file name.

Severity: Immediate Action

Explanation: The Idif2tdbm program encountered an error while processing a modify entry in a schema LDIF file. The Idif2tdbm program requires that each change clause group in a modify entry begin with a change indication line. There is no default value for this line.

System Action: The program ends because it cannot process the requested changes to the schema.

Administrator Response: Use the information in the message to locate the entry in the LDIF file and add the appropriate change indicator line. Then try the program again.

GLD3078A

Incorrect syntax is detected in an entry at line line_number of LDIF file file name.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program encountered an error while processing a modify entry in a schema LDIF file. Every line in an entry must follow the required syntax, which is type: value or type:: value.

System Action: The program ends because it cannot process the requested changes to the schema.

Administrator Response: Use the information in the message to locate the entry in the LDIF file and correct the syntax of the line. Then try the program again.

GLD3079A A change type that is not supported is found in an LDIF modify entry at line

line number of LDIF file file name.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program encountered an error while processing a modify entry in a schema LDIF file. The entry contains a change type line specifying a change type other than modify. The schema entry can only be modified.

System Action: The program ends because it cannot process the requested changes to the schema.

Administrator Response: Use the information in the message to locate the entry in the schema LDIF file and remove the entire change record. Then try the program again.

GLD3082A

The following option is not supported: option.

Severity: Immediate Action

Explanation: The program was invoked with an option that it does not support. The option that is not valid is displayed in the message.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again with the proper options.

GLD3083A

The DN value specified for the -s option contains characters that are not valid.

Severity: Immediate Action

Explanation: The **tdbm2ldif** program was invoked with a subtree DN value for the -s option that contains some characters that could not be converted or normalized.

System Action: The program ends.

Administrator Response: Ensure that the subtree DN value contains only valid characters. Then try the program again.

GLD3084A

No TDBM database suffix contains DN DN.

Severity: Immediate Action

Explanation: The program cannot find a TDBM database that includes a suffix that can contain the DN displayed in the message. Either no database suffix contains the DN or the database containing the DN is not a TDBM database.

For the Idif2tdbm program, this DN is either the DN of the first entry in the first non-empty schema LDIF file (if the -s or -v option is specified) or the DN of the first

entry of the first non-empty LDIF input file.

For the tdbm2ldif program, this is the subtree DN value specified for the **-s** option.

System Action: The program ends.

Administrator Response: Ensure that the configuration file used by the program includes the TDBM database that contains the DN displayed in the message and that this database is correctly configured. Also check that the syntax of the DN value is valid. Then try the program again.

GLD3085A

The tdbm2ldif program found more than one TDBM database section. Either use the -s or -n option to specify which TDBM section to process or remove all but one of the TDBM sections from the configuration file.

Severity: Immediate Action

Explanation: When the tdbm2ldif program is invoked without specifying the -s or the -n option, the program fails if there is more than one TDBM database section in the configuration file because it cannot determine which TDBM database to process.

System Action: The program ends.

Administrator Response: The tdbm2ldif program provides two options that are used to specify which one of the TDBM database sections in the configuration file to process. These options cannot both be specified at the same time.

- The -s option specifies a subtree whose entries are to be unloaded. The tdbm2ldif program selects the TDBM database section that contains this subtree from the configuration file.
- The -n option indicates the name of a TDBM database section whose entries are to be unloaded. The tdbm2ldif program selects the TDBM database section with this name from the configuration file.

Alternatively, you can modify the configuration file and remove all the TDBM database sections except for the one you want to process.

Then try the program again.

GLD3087A There is no entry in the LDAP directory for DN DN.

Severity: Immediate Action

Explanation: The DN displayed in the message has no entry in the LDAP directory. For the tdbm2ldif program, the DN is the subtree DN value specified for the **-s** option.

System Action: The program ends.

Administrator Response: For the tdbm2ldif program, either change the subtree DN value to be the DN for an

existing entry in the LDAP directory or remove the -s option. Then try the program again.

GLD3088A

SQL call returned unexpected return code return code on call: SQL call.

Severity: Immediate Action

Explanation: The program cannot continue because an unexpected return code was received from an SQL call. The return code and the SQL call are displayed in the message.

System Action: The program ends.

Administrator Response: Use the information in the message to fix the problem. Then try the program again.

GLD3089A

Cannot write record to output file file_name. Reason is: reason_text.

Severity: Immediate Action

Explanation: The program was not able to write a record to an output file. The output file can be a physical file or stdout (standard out). The reason text returned by the system is displayed in the message.

System Action: The program ends.

Administrator Response: Use the information in the message to fix the problem. Ensure that the directory or dataset name displayed in the message exists and that the file can be written. Also check that the file system or dataset is not full. Then try the program again.

GLD3090A

The schema LDIF file, file_name, contains an entry whose DN is not cn=schema, suffix where suffix is a suffix for this backend.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program was invoked with a schema LDIF file that cannot be processed because it contains an entry that is not a schema entry or that does not include a sufffix that belongs to this TDBM backend.

System Action: The program ends.

Administrator Response: Ensure that all the entries in the schema LDIF file specified in the message are schema entries. A schema entry DN must begin with cn=schema and the rest of the DN must be a suffix for this TDBM backend. The TDBM backend in use is the one which contains the suffix included in the DN of the first entry in the first schema file.

GLD3091A An internal exception occurred, rc = return_code. Exception text is:

exception_text.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program. The return code from the routine and the exception text are displayed in the message.

System Action: The program ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3092E Attention: an input file, file_name, has been changed since it was last used.

Severity: Eventual Action

Explanation: This invocation of **Idif2tdbm** uses the same output datasets (specified using the **-o** option) as a previous invocation, hence is seen as a continuation of processing of the earlier invocation. The **Idif2tdbm** program has detected that one of the input files specified in the **-s**, **-v**, **-i**, or **-e** option on this invocation has been modified since the earlier invocation. Since LDIF entries that were checked or prepared using the earlier version of this file may no longer be valid for the new contents of this file, this warning message is issued. The time (in seconds since the epoch) of when the earlier invocation of **Idif2tdbm** was run is stored in the TIME record in the status file, *hlq*.BULKLOAD.JCL(STATUS), where *hlq* is the value

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is 1 (yes), processing continues. If the response is 0 (no), or any character other than 1, processing ends.

specified in the **-o** option in the command.

Administrator Response: Determine whether processing should continue using the changes to the file specified in the message. If so, respond 1 (yes) to the prompt to allow processing to continue. Otherwise respond 0 (no), or any character except 1, to stop processing. Then try the program again, changing the requested processing steps to redo the prior steps with the modified input file. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3093E Attention: additional checking suspended for out-of-date files.

Severity: Eventual Action

Explanation: This invocation of **Idif2tdbm** uses the same output datasets (specified using the **-o** option) as a previous invocation, hence is seen as a continuation of processing of the earlier invocation. When **Idif2tdbm** detects that an LDIF input file has been modified, it issues a warning message. To prevent issuing too many of these messages, the program suspends additional file checking and issues this message when it reaches a pre-set limit to the number of messages.

System Action: The program continues to check for other warning conditions, then issues a message prompting whether to continue despite all the warning messages. If the response is 1 (yes), processing continues. If the response is 0 (no), or any character other than 1, processing ends.

Administrator Response: The previous warning message indicates the last LDIF input file that was checked. You should check the remaining LDIF input files to ensure that they are not out of date or that any changes to the files are acceptable. If so, respond 1 (yes) to the prompt to allow processing to continue. Otherwise respond 0 (no), or any character except 1, to stop processing. Then try the program again, changing the requested processing steps to redo the prior steps with the modified input files. You can use the -a option to provide a response to the prompt as part of the command invocation.

GLD3096A Idif2tdbm cannot add a child entry under the schema directory node.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program has detected an entry whose parent is the schema entry in the directory. The schema node cannot have any children. The child entry cannot be added to the directory.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Use the information in the accompanying message to locate the child entry in an LDIF input file and remove it from the file. Then try the program again.

GLD3097A Idif2tdbm cannot add a child entry without a parent.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program cannot add the LDIF entry to the directory because it cannot locate the parent of the entry. The parent either must already be in the database, or must be an entry before this entry in

this LDIF input file or in an LDIF input file processed before this file.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Use the information in the accompanying message to locate the child entry in an LDIF input file. Either remove the child entry or add an entry for the parent before the child entry. Then try the program again.

GLD3098A Schema entry found outside a schema input file.

Severity: Immediate Action

Explanation: The Idif2tdbm program provides the -s and -v options for specifying one or more schema input files, which contain entries to modify the schema. The program has found a schema entry in a different LDIF input file. The schema entry cannot be added to the directory. A schema entry DN begins with cn=schema and the rest of the DN is a suffix for the TDBM backend.

System Action: The program ends.

Administrator Response: Use the information in the accompanying message to locate the entry in an LDIF input file. Move the entry to a schema input file. Then try the program again. You can also use the Idapmodify utility or an LDAP client to process the schema entry.

GLD3099A Idif2tdbm cannot add a child entry under a referral object.

Severity: Immediate Action

Explanation: The Idif2tdbm program has detected an entry whose parent is a referral object, so the program cannot add the entry to the directory. The entry must instead be added to the directory that contains the actual entry which is referenced by the referral object.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Use the information in the accompanying message to locate the child entry in an LDIF input file and remove the child entry from the file. Then try the program again.

GLD3100A

Idif2tdbm cannot add a child entry under this parent because the parent has an incomplete entry in the database.

Severity: Immediate Action

Explanation: The Idif2tdbm program has detected an entry whose parent is in the database but the program cannot retrieve all the needed information about the

parent from the database. The incomplete information is either the ACL attributes of the parent or the set of ancestors of the parent. The child entry cannot be added to the directory.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Use the information in the accompanying message to locate the entry in an LDIF input file. Ensure that the database contains complete ACL attributes and ancestor information for the parent of this entry. Then try the program again.

GLD3101A Idif2tdbm cannot add a duplicate entry.

Severity: Immediate Action

Explanation: The Idif2tdbm program has detected an entry which is a duplicate. Another entry with the same DN is either already in the database, or is specified in an entry before this entry in this LDIF input file or in an LDIF input file processed before this file. The duplicate entry cannot be added to the directory.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Use the information in the accompanying message to locate the duplicate entry in an LDIF input file. If this entry is a duplicate of a previous entry in an LDIF input file, remove one of these entries. If the entry is a duplicate of an entry in the directory, remove the entry from the LDIF input file. Then try the program again.

GLD3102A The entry has failed schema check.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program has detected an entry whose attributes violate the current schema definition. The current schema is the schema in the directory plus any modifications contained in the schema LDIF files specified with the -s or -v options on the Idif2tdbm command. The entry cannot be added to the directory.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Use the information in the accompanying messages to locate the entry in an LDIF file and determine the problem. If the current schema needs to be modified, create an LDIF file containing the modification and specify the file using the -s or -v option on the **Idif2tdbm** command. Then try the program again.

GLD3104A Unable to open file file_name. Return code from fopen is errno; reason is: reason text

Severity: Immediate Action

Explanation: The file named in the message cannot be opened in the required way: read for an input file, write for an output file. The error code and reason text from **fopen** are displayed in the message.

System Action: Usually, the program ends. In some cases, the program can continue limited processing.

Administrator Response: For an input file, ensure that the file exists and can be opened for read. For an output file, check that the directory or dataset containing the file exists and that the file can be written to that directory or dataset. Then try the program again.

GLD3105A A creator DN must be specified for the entries to be loaded.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program requires a DN that the program can use to identify the creator of the entries that will be added to the directory, if those entries do not already include a creator attribute. This DN can be specified on the **-b** option of the **Idif2tdbm** command. If the **-b** option is not specified, the program uses the value of the adminDN record in the LDAP configuration file. The same value is also used for the modifier of the entries, if those entries do not already include a modifier attribute.

System Action: The program ends.

Administrator Response: Either add the -b option to the Idif2tdbm command or add the adminDN record to the LDAP configuration file. Then try the program again.

GLD3106A The schema has changed since the load files were prepared.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program has terminated load processing because it has detected that the data in the load files may be out of date. The current schema in the directory was modified after the load files were prepared, thus the data in the load files may not be valid for the modified schema. The **Idif2tdbm** program compares the time of when the load files were prepared (saved in the status file) with the time of when the schema was last modified (saved in the DIR_CACHE table in the database) to determine if the load files may not be usable. The time when the load files were prepared (in seconds since the epoch) is stored in the TIME record in the status file, *hlq*.BULKLOAD.JCL(STATUS), where *hlq* is the value specified in the **-o** option in the

System Action: The program ends.

Administrator Response: Run Idif2tdbm with the -c and -p options to check and prepare a new set of load files using the current schema. Then invoke the program again to load the data. If you are absolutely sure that the data in the load files is valid for the current schema, you can 'trick' Idif2tdbm into loading the data by resetting the value of the TIME record in the status file to the present time. Then try the program again.

GLD3107A

An internal exception occurred creating cache manager, rc = return_code. Exception text is: exception text

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program. The return code from the routine and the exception text are displayed in the message.

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3108A

An internal exception occurred creating attribute cache, rc = return_code. Exception text is: exception_text

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program. The return code from the routine and the exception text are displayed in the message.

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3109A Initialization of database tables failed, rc = return_code.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine while attempting to initialize database tables. The return code from the routine indicates the type of SQL error encountered.

System Action: The program continues. The backend

which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3110A **Error formatting data for** *table_name* tables, rc = return code.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program to format data before updating the database. The return code from the routine is displayed in the message

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3111A SQL Error updating table_name tables, rc = return_code.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine while attempting to update database tables. The return code from the routine indicates the type of SQL error encountered.

System Action: The program continues. The backend which encountered the error ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3112A

An internal exception occurred creating schema cache, rc = return_code. Exception text is: exception_text.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program. The return code from the routine and the exception text are displayed in the message.

System Action: The program continues. The backend

which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3113A An error occurred creating schema cache from database entry cn=schema.

Severity: Immediate Action

Explanation: An error occurred creating an internal cache of the schema entry read from the database.

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3114A An unknown internal exception occurred during process processing.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program. The name of the routine is displayed in the message.

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD3115A The list file file name does not contain any LDIF input file names.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program requires LDIF input files to be specified on the command line. The -e option allows specifying a file that contains a list of LDIF input file names. If the -i option is not specified and the file specified with the **-e** option does not contain any LDIF input file names, the **ldif2tdbm** program cannot continue.

System Action: The program ends.

Administrator Response: Specify a new file with the -e option that contains valid LDIF input file names, add valid LDIF file names to the file specified in the

message, or specify LDIF input files with the **-i** option. Then try the program again.

GLD3116A The file file_name is missing a required record: keyword.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program requires that one SSID, one HLQ, and at least one JOBCARD record be specified in system file. The record specified in the message is missing from the system file.

System Action: The program ends.

Administrator Response: Add the required record to the system. The try the program again.

GLD3118A Record without value found in file file name on line line number.

Severity: Immediate Action

Explanation: The **Idif2tdbm** program uses a system and status file for required information. The format of the records in these files is: a name, followed by any number of blanks, followed by a value. The only valid names for the system file are HLQ, SSID, and JOBCARD. The only valid names for the status file are STATUS, TIME, LDIFFILE, and SCHEMA.

System Action: The program ends.

Administrator Response: Either add a value to the record or remove the record from the file. Then try the program again.

GLD3119I A DB2 deadlock/timeout has been detected and the transaction is being

retried internally.

Severity: Information

Explanation: The TDBM backend detected a DB2 deadlock/timeout. The transaction is being retried.

System Action: The program continues.

Operator Response: None.

Administrator Response: None.

GLD3120A Idif2tdbm could not submit JCL file name.

Severity: Immediate Action

Explanation: The load step of the **Idif2tdbm** program failed when it attempted to submit the JCL file specified in the message to the JES reader. The reason for the failure is explained in a previous message.

System Action: The program ends.

Administrator Response: Use the information in any error messages issued by the command to fix the problem. **Note:** Do not run the **Idif2tdbm** program again

because this can add duplicate data to the database. Instead, use the information in the description of the **Idif2tdbm** command in *z/OS Security Server LDAP Server Administration and Use Guide* to determine how to proceed.

GLD3121A The record_name record in the system file file_name is not valid.

Severity: Immediate Action

Explanation: The Idif2tdbm program uses the information in the system file to create the JCL used to load the database. The program detected a problem with a record in the status file. The problem is either that the format of the value on the first JOBCARD record is not valid or that the value on one of the records is too long. The value of the first JOBCARD record must begin with the job name, which consists of 2 slashes directly followed by up to 8 characters (for a total length of up to 10 characters). For usage in the JCL, the maximum length is 71 for a JOBCARD record value, 55 for an SSID record value, and 36 for an HLQ record value. Note: DB2 may have different length constraints for some of these values. The record values must meet both the DB2 and Idif2tdbm length limitations.

System Action: The program continues to process the check step if it has been requested; otherwise, the program ends.

Administrator Response: Correct any problems with the syntax of the first JOBCARD record value or with the lengths of any of the record values in the system file. Then try the program again.

GLD3122A The __passwd function failed; not loaded from a program controlled library.

Severity: Immediate Action

Explanation: The LDAP server with a TDBM backend needs to be loaded from a program controlled dataset for the __passwd function to work.

System Action: The client request fails with an operations error. The program continues.

Operator Response: Ensure that the LDAP server and the TDBM backend are loaded from a program controlled dataset. Stop and start the server after checking the dataset.

GLD3123A TDBM backend __passwd internal error. Program continues.

Severity: Immediate Action

Explanation: The TDBM backend received an unexpected error from the __passwd() function while processing a bind request.

System Action: The client request fails with an

operations error. The program continues.

continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: If the problem persists, contact the service representative.

GLD3124E A Kerberos princSubtree does not exist in the directory. Program

Severity: Eventual Action

Explanation: The princSubtree that exists in the REALM object was not found in the directory. The server will continue to operate but unexpected results may occur.

System Action: The program continues but unexpected results may occur.

Administrator Response: Either the princSubtree value was incorrect or the entry needs to be added to the directory. Correct the error and try again.

GLD3127A Unable to establish initial caches.

Severity: Immediate Action

Explanation: The LDAP server TDBM backend encountered an error while trying to establish initial caches.

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: One reason for this error may be lack of memory. Ensure the server has enough memory and restart the server. If the problem persists, contact the service representative.

GLD3128A Error establishing suffixes.

Severity: Immediate Action

Explanation: The LDAP server TDBM backend encountered an error while normalizing one of the suffix values in the configuration file.

System Action: The program continues. The backend which encountered the error ends.

Operator Response: This error occurs if the server has run out of memory. Ensure the server has enough memory and restart the server.

GLD3129E attrOverflowSize out of bounds; using default value of 255.

Severity: Eventual Action

Explanation: The value specified for attr0verflowSize in the configuration file is out of bounds; therefore, the value defaults to 255.

System Action: The program continues.

Administrator Response: If a different

attr0verflowSize is desired, stop the server, correct the configuration file and restart the server. The value can not be larger than the maximum integer supported on the system.

GLD3130E useNativeAuth value not valid; using default value of 'off'.

Severity: Eventual Action

Explanation: The value specified for useNativeAuth in the configuration file is not valid; therefore, the value defaults to off.

System Action: The program continues. Native authentication support will not be activated in the server.

Administrator Response: If a different useNativeAuth value is desired, stop the server, correct the configuration file and restart the server. Valid values are on, yes, off and no.

GLD3131E nativeUpdateAllowed value not valid; using default value of 'no'.

Severity: Eventual Action

Explanation: The value specified for nativeUpdateAllowed in the configuration file is not valid; therefore, the value defaults to no.

System Action: The program continues. Native password update support will not be activated in the server.

Administrator Response: If a different nativeUpdateAllowed value is desired, stop the server, correct the configuration file and restart the server. Valid values are **yes** and **no**.

The useNativeAuth configuration GLD31321 option all/selected has been enabled.

Severity: Information

Explanation: Native authentication is enabled for this

backend.

System Action: None.

Administrator Response: None.

GLD3133A The version (dbversion) of the TDBM database is not supported by the current program level.

Severity: Immediate Action

Explanation: The program cannot interact with the TDBM database due to a level mismatch. This could be due to running an earlier program version after the TDBM database version has been updated.

System Action: The program continues. The backend which encountered the error ends.

Administrator Response: Ensure that the correct program level is being loaded. Refer to the migration section of *z/OS Security Server LDAP Server Administration and Use* to determine the proper TDBM database version for the program level being used.

GLD3134E No -o parameter supplied, writing to standard output.

Severity: Eventual Action

Explanation: The **tdbm2ldif** program requires a location where it can place the LDIF-formatted output. If the **-o** parameter is not supplied, **tdbm2ldif** will write to standard output.

System Action: The program continues.

Administrator Response: No action may be necessary. If the output from tdbm2ldif was intended for a file, run the program again, using the -o parameter to indicate where the program should place its output. The value of the option must be the path name of a file or a dataset for tdbm2ldif.

GLD3135I Grant/Deny ACL support is enabled

below suffixes: suffix-list.

Severity: Information

Explanation: Grant/Deny style ACLs along with attribute-level permission settings are supported below the specified suffixes.

System Action: The program continues.

Operator Response: None.

Administrator Response: None.

·

GLD3136l Grant/Deny ACL support is not enabled below suffixes: suffix-list.

Severity: Information

Explanation: Grant/Deny style ACLs along with attribute-level permission settings are not enabled below the specified suffixes. The database version indicates that there may be systems accessing the database that do not support the additional formats.

System Action: The program continues.

Operator Response: None.

Administrator Response: Refer to the migration section of *z/OS Security Server LDAP Server Administration and Use* for information on migrating the database. The database should only be migrated after all servers accessing the database can support the new database formats.

GLD3137I Using TDBM backend named 'name'.

Severity: Information

Explanation: The program is processing the TDBM backend named in the message. The name corresponds to the name specified on the database record marking the beginning of the information for this TDBM backend in the configuration file.

System Action: The program continues.

Operator Response: None.

Administrator Response: You can use the name to locate the database section for this TDBM backend in the configuration file and check that this is the backend that you want to process.

GLD3138I Using TDBM backend containing suffix 'suffix'.

Severity: Information

Explanation: The program is processing the TDBM backend containing the suffix displayed in the message. This suffix corresponds to the value of the first suffix record for this TDBM backend in the configuration file.

System Action: The program continues.

Operator Response: None.

Administrator Response: You can use the suffix to locate the database section for this TDBM backend in the configuration file and check that this is the backend that you wanted to process.

GLD3139A There is no TDBM backend with name 'name'.

Severity: Immediate Action

Explanation: The program cannot find a TDBM backend with the name displayed in the message. For the **tdbm2ldif** program, this name is the value specified for the **-n** option.

System Action: The program ends.

Administrator Response: Ensure that the configuration file used by the program includes a database record that contains the name displayed in the message and specifies tdbm for the database type. Also make sure that the backend is correctly configured. Then try the program again.

GLD3140A

TDBM failed during initialization. One of the following is incorrect: database userid='dbuserid'; Table name='tablename'; Table column='tablecolumn'.

Severity: Immediate Action

Explanation: The program cannot find some global

information for the TDBM backend. Either the database userid, table name, or table column is wrong.

System Action: The program ends.

Administrator Response: Ensure that the dbuserid is correct in the configuration file. If the dbuserid is correct, then ensure that there are DB2 tables created for this TDBM backend.

GLD31411

End Successful Modify DN operation: newrdn =>

new_Relative_Distinguished_Name; deleteoldrdn => deleteOldRdn_flag; dn => Distinguished_Name

Severity: Information

Explanation: The LDAP server successfully completed a Modify DN operation. This message displays the new RDN, deleteOldRdn flag and target DN values, respectively, in the operation request. This information is provided to track the changes requested by Modify DN operations, because they may potentially affect many entries in the directory.

System Action: The program continues.

Operator Response: None. Administrator Response: None.

GLD3142I

End Successful Modify DN operation: newSuperior =>

new_Superior_Distinguished_Name; dn => Distinguished_Name

Severity: Information

Explanation: The LDAP server successfully completed

a Modify DN operation. This message displays the new Superior DN and target DN respectively, in the operation request. This information is provided to track the changes requested by Modify DN operations, because they may potentially affect many entries in the directory.

System Action: The program continues.

Operator Response: None. Administrator Response: None.

GLD3143I

End Successful Modify DN operation: control type => Control_Type; control criticality => Control Criticality; control value (hexadecimal) => x'Control_Value'; dn => Distinguished_Name

Severity: Information

Explanation: The LDAP server successfully completed a Modify DN operation. This message displays the contents of any controls and the target DN submitted in the operation request, including control type, control criticality, and control value (shown in hexadecimal format) for each control. This information is provided to track the changes requested by Modify DN operations, because they may potentially affect many entries in the directory.

System Action: The program continues.

Operator Response: None. Administrator Response: None.

LDAP server product messages (4000)

GLD4001A

Requested message number max_number beyond bounds of internal table.

Severity: Immediate Action

Explanation: The message specified by the message number cannot be displayed because it is not in the internal message table and was not found in any catalog.

System Action: The program continues.

Operator Response: Contact the service representative. An internal error has occurred.

GLD4002E Cannot open message catalog file file name.

Severity: Eventual Action

Explanation: The program was unable to locate the specified message catalog. Possible reasons include:

message catalogs installed incorrectly or, LANG or **NLSPATH** environment variables may not be set or may be set incorrectly.

System Action: The program continues.

Operator Response: Verify the full path name to the message catalogs. Verify that the LANG and NLSPATH variables have the correct values to reach this path. Ensure the message catalogs are installed correctly and have the correct permissions. If the problem persists, contact your service representative.

GLD4003A Memory allocation failed; cannot continue. Program terminated.

Severity: Immediate Action

Explanation: An LDAP facility was unable to allocate the necessary memory to continue processing. The program is ending.

System Action: The program ends.

Operator Response: Ensure that the program has sufficient memory and try again. If the problem persists, contact the service representative.

GLD4004A Unable to write error message to console. Return code=return code.

Severity: Immediate Action

Explanation: An LDAP program received an error from system routine **__console()** when attempting to write a message to the operator's console.

System Action: The program continues.

Operator Response: Contact the service

representative.

GLD4005E Environment variable file not found. Environment variables not set. Continuing.

Severity: Eventual Action

Explanation: The LDAP program was unable to find the specified environment variable file. Program continues with environment variables unset.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Verify the existence and location of the environment variable file and restart the server.

GLD4006E Environment variable file cannot be opened. Environment variables not set.

Continuing.

Severity: Eventual Action

Explanation: The LDAP program was unable to open the environment variable file. Program continues with environment variables unset.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Verify that the file system is operating correctly and restart the server.

GLD4007E Environment variable file contents in error. Environment variables not set. Continuing.

Severity: Eventual Action

Explanation: The LDAP program encountered an incorrect line in its environment variable file. Program continues with environment variables unset.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Verify the contents of the environment variable file and restart the server.

GLD4008E Environment variables not set because error encountered. Continuing.

Severity: Eventual Action

Explanation: The LDAP program encountered an unexpected error when processing its environment variable file. Program continues with environment variables unset.

System Action: The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD4009A The initialization of ServerGlobal structure failed.

Severity: Immediate Action

Explanation: The LDAP program encountered an error when establishing the ServerGlobal data structure. The program cannot continue without this structure.

System Action: The program ends.

Administrator Response: Ensure adequate storage for the program and try again. If the problem persists, contact the service representative.

GLD4010A The initialization of ConfigInfo structure failed.

Severity: Immediate Action

Explanation: The LDAP program encountered an error when establishing the ConfigInfo data structure. The program cannot continue without this structure.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Ensure adequate storage for the program and try again. If the problem persists, contact the service representative.

GLD4012A The initialization of ReplInfo structure failed.

Severity: Immediate Action

Explanation: The LDAP program encountered an error when establishing the Replinfo data structure. The

program cannot continue without this structure.

System Action: The program ends.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: Ensure adequate storage for the program and try again. If the problem persists, contact the service representative.

GLD4013A Unable to locate specified catalog.

Severity: Immediate Action

Explanation: The LDAP program encountered an error when attempting to issue a message. The program will continue but may not issue appropriate messages.

System Action: The program continues. Operator Response: Contact the service

representative.

GLD4014A Unable to establish message support.

Severity: Immediate Action

Explanation: The LDAP program encountered an error when attempting to establish message support.

System Action: The program ends.

Operator Response: Contact the service

representative.

GLD4015A An error occurred when reading the initial schema.

Severity: Immediate Action

Explanation: The LDAP server or a server backend encountered an error while trying to read the initial or minimum schema.

System Action: If the problem occurs while processing a backend, the server continues to process any other backends. Otherwise, the server ends.

Administrator Response: If the initial schema is being read from a file, make sure that the file exists and can be read. Additional information is provided using debug messages when the debug level is set to LDAP_DEBUG_ERROR. If the problem persists, contact the service representative.

GLD4016A Unable to establish initial schema. Return code is: return_code; reason is:

reason_text

Severity: Immediate Action

Explanation: The LDAP server or a server backend encountered an error while trying to establish the initial or minimum schema.

System Action: If the problem occurs while

processing a backend, the server continues to process any other backends. Otherwise, the server ends.

Administrator Response: Use the LDAP return code and reason text displayed in the message to resolve the problem. Then start the server again. If the problem persists, contact the service representative.

GLD4017A Error in internal routine routine_name, rc = return_code.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program. The name of the routine and its return code are displayed in the message.

System Action: If the program is the LDAP server and the problem occurs while processing a backend, the server continues to process any other backends. Otherwise, the program ends.

Administrator Response: Use the information in any error or debug messages issued by the routine to fix the problem. Then try the program again. If the problem persists, contact the service representative.

GLD4018A Option 'option_name' is specified more than once with different values.

Severity: Immediate Action

Explanation: The invocation of this program includes an option that is specified more than once with a different value. The program cannot determine which value to use for the option. The option is displayed in the message.

System Action: If the program is the LDAP server and the problem occurs while processing a backend, the server continues to process any other backends. Otherwise, the program ends.

Administrator Response: Ensure that the option is specified only once and has the desired value. Then try the program again.

GLD4019A The following options cannot both be specified: 'option1' and 'option2'.

Severity: Immediate Action

Explanation: The program was invoked with two options that are mutually exclusive. You cannot specify both of these options.

System Action: The program ends.

Administrator Response: Refer to the usage information for the program and try the program again with the proper options.

GLD4020I The following LDAP message is truncated at the console.

Severity: Information

Explanation: The next console message from LDAP (prefix GLD) is longer than the console message limit. See the LDAP server log to view the complete message.

System Action: The program continues.

Operator Response: None.

Administrator Response: See the server log for the

complete message.

SDBM messages (5000)

GLD5001A Memory allocation failed; cannot continue. Program terminated.

Severity: Immediate Action

Explanation: The SDBM backend was unable to allocate the necessary memory to continue processing. The program is ending.

System Action: The program ends.

Operator Response: Ensure that the LDAP server has sufficient memory and try again. If the problem persists, contact the service representative.

GLD5002A The __passwd() function failed; not loaded from a program controlled library.

Severity: Immediate Action

Explanation: The LDAP server with an SDBM backend needs to be loaded from a program controlled dataset for the __passwd() function to work.

System Action: The client request fails with an operations error. The program continues.

Operator Response: Ensure that the LDAP server and the SDBM backend are loaded from a program controlled dataset. Stop and start the server after checking the dataset.

GLD5003A SDBM backend __passwd() internal error. Program continues.

Severity: Immediate Action

Explanation: The SDBM backend received an unexpected error from the __passwd() function while processing a bind request.

System Action: The client request fails with an operations error. The program continues.

Operator Response: Contact the LDAP Administrator or see the Administrator Response.

Administrator Response: If the problem persists, contact the service representative.

GLD5004A keyword parameter is missing or has no value in the configuration file.

Severity: Immediate Action

Explanation: The SDBM backend for the LDAP server cannot be configured because the configuration file is missing a required parameter or the parameter has no value.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Correct the configuration file and try again.

GLD5005A RDN in suffix 'suffix' is not valid.

Severity: Immediate Action

Explanation: The value for the suffix parameter is not valid in the SDBM backend section of the configuration file. The specified Relative Distinguished Name (RDN) is not correct.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Correct the suffix value in the configuration file. Make sure that the RDN consists of an attribute type and value, separated by an equal sign.

GLD5006A Only one suffix can be specified for the SDBM backend in the configuration file.

Severity: Immediate Action

Explanation: The SDBM backend for the LDAP server cannot be configured because the configuration file contains multiple suffix lines for SDBM. There can only be one one SDBM backend section in the configuration file and it must contain only one suffix line.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Correct the configuration

file and try again.

Idapenf messages (7000)

GLD7001I Usage: Idapcnf -i profile_file

Severity: Information

Explanation: The correct syntax of the **Idapcnf**

command is: Idapcnf -i profile file

System Action: The program ends.

Operator Response: Invoke the Idapcnf command again with the appropriate options.

Administrator Response: None.

GLD7002I Generating output_file

Severity: Information

Explanation: Status message which indicates the generation of an output file by the Idapcnf utility.

System Action: The program continues.

Administrator Response: None.

GLD7003I Finished generating output_file.

Severity: Information

Explanation: Status message which indicates the completed generation of an output file by the Idapcnf

utility.

System Action: The program continues.

Administrator Response: None.

GLD7004A Errors found. Correct errors and try again.

Severity: Immediate Action

Explanation: The **Idapcnf** utility detected errors.

System Action: The program ends.

Administrator Response: Use the information in any error messages issued by the utility to fix the problem.

Then try the command again.

GLD7005I The utility is finished checking for errors.

Severity: Information

Explanation: Status message which indicates that

Idapcnf has completed error checking.

System Action: The program continues.

Administrator Response: None.

GLD7006A statement_name is missing in the input

file input_file.

Severity: Immediate Action

Explanation: The **Idapcnf** utility has detected a

missing statement in the input file.

System Action: The program continues.

Administrator Response: Correct the missing

statement in the input file.

GLD7007A Value statement_name is greater than

length characters in the input file

input file.

Severity: Immediate Action

Explanation: The **Idapcnf** utility has detected that the statement value is greater than the specified limit in the

input file.

System Action: The program continues.

Administrator Response: If the value should be greater than the limit specified, see the Usage Notes for the Idapcnf utility. Otherwise correct the value in the

input file.

GLD7008A Value for statement_name contains nonprintable characters in the input

file input_file.

Severity: Immediate Action

Explanation: The Idapcnf utility has detected that a specified value in the input file contains nonprintable

characters.

System Action: The program continues.

Administrator Response: Correct the statement value

in the input file.

GLD7009E

Attention: The output data set data_set_name has been previously used. Do you wish to overwrite the existing members of this data set?

(y/n) [n]

Severity: Eventual Action

Explanation: The output data set specified on Idapcnf invocation already contains outputs from the command. This is prompting the user to specify if they wish to overwrite existing members in the output data set. If the output data set is currently being used for an LDAP server, a different output data set should be used for

this invocation of the Idapcnf utility.

System Action: The program continues or ends

based upon input.

Administrator Response: Specify y or n.

GLD7010A The temporary directory

(TEMPORARY_DIR statement in the input file input_file) specified (directory_name) does not exist.

Severity: Immediate Action

Explanation: The **Idapcnf** utility has detected that the directory specified on the TEMPORARY_DIR statement

does not exist.

System Action: The program ends.

Administrator Response: Correct the temporary directory statement (TEMPORARY_DIR) value in the

input file.

GLD7011I Exiting with return code return code.

Severity: Information

Explanation: The **Idapcnf** utility is exiting.

System Action: The program ends.

Administrator Response: If the return code is nonzero, look for previous error messages.

GLD7012I Terminating upon user request.

Severity: Information

Explanation: The Idapcnf utility has been terminated

by the user.

System Action: The program ends.

Administrator Response: None.

Extended operations messages (9000)

GLD9001A An error occurred during backend_type backend initialization, rc = return_code.

Severity: Immediate Action

Explanation: An error occurred during backend

initialization.

System Action: The program continues. Backends that configure successfully will be available. If no backends configure successfully, an additional message will appear and the program will end.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD9002A An unknown internal exception occurred during an extended operation.

Severity: Immediate Action

Explanation: An unexpected error occurred in an

internal routine used by the program.

System Action: The program continues.

Administrator Response: Contact the service

representative.

GLD9003A An internal exception occurred during an extended operation. Return code is

errno; error string is: error_string.

Severity: Immediate Action

Explanation: An unexpected error occurred in an internal routine used by the program.

System Action: The program continues. The request fails.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD9004E The extended operations backend encountered an error while attempting LDAP client requests. Return code is

ldap_errno; error string is:
ldap_errstring.

Severity: Eventual Action

Explanation: To satisfy some extended operations, the extended operations backend uses the LDAP client APIs. While handling such a request, the extended operations backend encountered an error.

System Action: The program continues. The request fails.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD9005E The extended operations backend

encountered an unrecognizable control. Return code is *Idap_errno*; error string is: *error_string*.

Severity: Eventual Action

Explanation: During an extended operations request,

the extended operations backend encountered an unrecognizable control.

System Action: The program continues. The request fails.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD9006E

The LDAP program encountered an unrecognizable extended operation request. Return code is Idap_errno; error string is: error_string.

Severity: Eventual Action

Explanation: The LDAP program encountered an unrecognizable extended operations request.

System Action: The program continues. The request

fails.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD9007E

The extended operation requires a critical control. Error string is:

error_string.

Severity: Eventual Action

Explanation: The LDAP program did not receive a required critical control with an extended operations request.

System Action: The program continues. The request fails.

Administrator Response: Use the information in the message to resolve the problem. For more detailed error information, issue the command again, specifying the debug error option and use the debug information to try to resolve the problem. If the problem persists, contact the service representative.

GLD9008A The extended operation backend could

not allocate memory. Error string is:

error_string.

Severity: Immediate Action

Explanation: The LDAP program could not allocate

memory during its processing.

System Action: The program continues. The request

Operator Response: Contact the LDAP Administrator

or see the Administrator Response.

Administrator Response: Ensure adequate storage for the program and try again. If the problem persists, contact the service representative.

Part 4. Appendixes

Appendix A. LDAP server configuration file (slapd.conf)

This appendix shows the LDAP server configuration file, slapd.conf, located in /usr/lpp/ldap/etc.

```
*******************
# This file is shipped in code page IBM-1047 and must remain in
# code page IBM-1047.
 *****************
 *****************
# Licensed Materials - Property of IBM
 5694-A01
# (C) Copyright IBM Corp. 1997, 2001
# ******************
# Filename slapd.conf
# This is the main configuration file for the z/OS Security
# Server LDAP Server. Refer to publication:
# z/OS Security Server LDAP Server Administration and Usage
# Guide (SC24-5923)
# for details on specifying each configuration parameter in this
# The LDAP configuration is specified in sections. At the top
# of the file is the main configuration section into which
# global configuration parameters are specified. Items such as
# the maximum number of connections, number of threads,
# SSL support, and schema definitions that are common to multiple
# portions of the server are specified in the top section.
# After the top section come database-specific configuration sections.
# These sections setup configuration parameters that are specific to
# a particular backend.
# The configuration switches from top section to database-specific
# section when a database configuration line is encountered. All
# configuration parameters after a database configuration line are
# accepted as database-specific configuration parameters.
# When another database configuration line is encountered, all
 configuration parameters are now specific to the database specified
 on the database configuration parameters.
 Basic format of the LDAP configuration file:
# <top section>
# database sdbm GLDBSDBM
# <SDBM database-specific configuration section>
# database tdbm GLDBTDBM
# <TDBM database-specific configuration section>
# database rdbm GLDBRDBM
```

```
<RDBM database-specific configuration section>
 #-----
\# Within a section, different portions of the LDAP configuration
# file can be broken up into separate files. These files are included
# into the configuration using the include configuration parameter.
# The entire contents of the file referenced by the include parameter
# are treated as included where the include statement was encountered.
# Separate include statements are useful for isolating the RDBM-style schema
# definitions from the rest of the LDAP server configuration or for separating
# different database-specific options from the main configuration file.
\# Configuration lines are continued on the following line when the following
# line contains a whitespace character (space or tab character) as the first
 character of the line. Thus, the following two configuration lines are
 equivalent:
# attribute attr1 cis attr1col 128 normal
 attribute attr1
    cis
    attr1col
    128
    normal
# This configuration file provides information on all configuration options
# which can be specified. Configuration options that are usually specified
# are uncommented in the file.
TOP SECTION
# -----
-----
# include <filename>
# Default Value: none
# Example:
   include /etc/ldap/schema.system.at
   include //'sys1.ldap.parms'
# NOTES:
   Definitions of schema (objectclass and attribute configuration settings)
   are usually separated out into separately included files.
include    /etc/ldap/slapd.at.system
include    /etc/ldap/slapd.cb.at.conf
include
            /etc/ldap/slapd.at.conf
include
            /etc/ldap/slapd.at.racf
            /etc/ldap/slapd.oc.system
include
```

```
/etc/ldap/slapd.cb.oc.conf
include
include
              /etc/ldap/slapd.oc.conf
include
              /etc/ldap/slapd.oc.racf
 -----
# objectclass <name> requires <name[,name,...]> allows <name[,name,...]>
# Default Value: none
# Example:
   objectclass newPerson requires cn,sn allows userpassword,telephonNumber
   objectclass jobPerson requires cn,sn,jobTitle allows userpassword
 -----
# attribute <name [ name ... ] > <cis|ces|tel|bin|dn>
  <columnName> <size> <normal | sensitive | critical>
# Default Value: none
# Example:
  attribute jobTitle cis jobTitle 128 normal attribute jobDescription ces jobDesc 1000 normal
# verifySchema <on|off>
# Default Value: on
# Example:
   verifySchema on
verifySchema on
# -----
# sizeLimit <number>
# Default Value: 500
# Example:
  sizeLimit
               350
  sizelimit limits the number of entries that will be returned in
   a single search operation. The setting in this file limits
   the number of returned entries to 500.
sizeLimit 500
# -----
# timeLimit <seconds>
# Default Value: 3600
# Example:
   timeLimit
               60
```

```
# NOTES:
   timelimit limits the time that a connection can remain active
  with the server. The setting in this file allows connections
  to stay active for 1 hour (3600 seconds).
# -----
timeLimit 3600
 -----
# maxConnections <number>
# Default Value: 65535
# Example:
  maxConnections 200
# NOTES:
#
  The following settings define the server to allow 65535
   concurrent connections.
# -----
maxConnections 65535
 -----
# commThreads <number>
# Default Value: 10
# Example:
  commThreads 80
# NOTES:
  commThreads specifies the number of threads that are initialized
  and waiting in the communication thread pool. The communication
   thread pool processes requests from the clients and dispatches them
  to the appropriate backend.
# -----
commThreads 80
 -----
 pcThreads <number>
# Default Value: 10
# Example:
  pcThreads 5
# NOTES:
  pcThreads specifies the number of threads that are initialized
  and waiting in the PC callable thread pool. The PC callable
  thread pool processes requests from the z/OS LDAP SAF interface
   and dispatches them to the appropriate backend. This thread pool
   is only activated when a listen has been configured for the SAF
   interface.
# ______
#pcThreads 5
# -----
# idleConnectionTimeout
                       <seconds>
# Default Value: (indefinite if not specified)
# Example:
```

```
idleConnectionTimeout
                           60
# NOTES:
   idleConnectionTimeout specifies the amount of time in seconds
   that the LDAP server should wait on a particular client
   connection.
# ______
# idleConnectionTimeout 60
# listen <ldapURL>
# Default Value: ldap://:389
# Example:
  listen ldap://host.yourcompany.com:444
   listen ldap://host2.yourcompany.com
   listen ldaps://host.yourcompany.com:736
   listen ldaps://host2.yourcompany.com
  listen ldap://:777
  listen ldaps://:636
  listen ldap://:pc
# NOTES:
   The listen parameter replaces the use of port, securePort,
   and security parameters in the configuration file.
   More than one listen parameter may appear in the configuration
   file so that the LDAP server will listen on multiple
   addresses-port pairs for client requests.
   The default setting for the listen parameter is to listen
   to clients on the non-secure port of 389.
   Example explanations:
   ldap://host.yourcompany.com:444 - The LDAP server will listen on
          the non-secure port of 444 to client requests from
          host.yourcompany.com.
   ldap://host2.yourcompany.com - The LDAP server will listen on the
          non-secure default port of 389 to client requests from
          host2.yourcompany.com.
   ldaps://host.yourcompany.com:736 - The LDAP server will listen on
           the secure port of 736 to client requests from
           host.yourcompany.com.
   ldaps://host2.yourcompany.com - The LDAP server will listen on the
           default secure port of 636 to client requests from
           host2.yourcompany.com.
   ldap://:777 - The LDAP server will listen on the non-secure port
           of 777 to client requests from any hostname or IP address.
   ldaps://:636 - The LDAP server will listen on the secure port
           of 636 to client requests from any hostname or IP address.
   ldap://:pc - The LDAP server will listen for PC calls on the z/OS
               SAF interface.
listen ldap://:389
#listen ldaps://:636
# -----
```

```
# altServer <ldapURL>
# Default Value: none
# Example:
   altserver ldap://otherhost.yourcompany.com:444
   altserver ldap://otherhost2.yourcompany.com
# NOTES:
   More than one altserver setting may appear in the configuration
   if there are multiple alternate servers.
# altserver ldap://otherhost.yourcompany.com:444
# -----
# referral <1dapURL>
# Default Value: none
# Example:
  referral ldap://ldap.yourcompany.com
# NOTES:
  Use this option to refer clients to a different server which has
   access to other parts of the directory name-space.
# referral ldap://ldap.yourcompany.com
 adminDN <distinguishedname>
# Default Value: none
# Example:
   adminDN "cn=Admin, o=Your Company"
# NOTES:
   This configuration option MUST be specified.
  The adminDN option should be updated to contain a
   distinguished name within one of the suffixes defined below.
   This requires that an entry exist in the directory for this
   distinguished name and it will be used when evaluating an
   LDAP bind operation for the AdminDN.
# adminDN "cn=Admin, o=Your Company"
 _____
 adminPW <password>
# Default Value: none
# Example:
  adminPW secret
# NOTES:
  This configuration option is optional and NOT recommended for general
  use. It is recommended that this option be used only during the initial
   startup of the LDAP server only until an entry in the LDAP directory
   can be defined which will serve as the administrator.
```

```
sysplexgroupname <name>
# Default Value: none
# Example:
   sysplexgroupname LDAPGRP
# sysplexservername <name>
# Default Value: none
# Example:
   sysplexservername LDAPSVR1
# sslAuth <serverClientAuth|serverAuth>
# Default Value: serverAuth
# Example:
  sslAuth serverClientAuth
# sslAuth serverAuth
# sslCertificate <certificateLabel|none>
# Default Value: none
# Example:
   sslCertificate cert1
# sslCertificate none
# -----
# sslCipherSpecs <number>
# Default Value: 0
# Example:
   sslCipherSpecs 15104
# NOTES:
   The configuration option value is specified in decimal but represents
   a mask of bits which indicate what cipher specs to accept.
# sslCipherspecs 15104
# sslKeyRingFile <filename>
```

```
# Default Value: none
# Example:
   sslKeyRingFile /etc/ldap/key.kdb
  sslKeyRingFile LDAPRING
# NOTES:
  If you are using RACF keyring support, specify the RACF keyring name
   in this field and do not specify either sslKeyRingFilePW or
   sslKeyRingPWStashFile in the configuration.
# sslKeyRingFile LDAPRING
# -----
 sslKeyRingFilePW <password>
# Default Value: none
# Example:
  sslKeyRingFilePW password
# NOTES:
  Use of this option is NOT recommended. Instead, use either the
  sslKeyRingPWStashFile configuration setting or RACF keyrings to
   manage the server certificate for SSL.
 -----
 sslKeyRingPWStashFile <filename>
# Default Value: none
# Example:
   sslKeyRingPWStashFile /etc/ldap/key.sth
# -----
# validateincomingV2strings <on|yes|off|no>
# Default Value: yes
# Example:
  validateincomingV2strings yes
validateincomingV2strings yes
 -----
 sendV3stringsoverV2as <ISO8859-1|UTF-8>
# Default Value: UTF-8
# Example:
   sendV3stringsoverV2as UTF-8
sendV3stringsoverV2as UTF-8
```

```
supportKrb5 <yes no>
 Default Value: no
 Example:
   supportKrb5 yes
# NOTES:
   If you are going to authenticate via Kerberos binds then this option
   should be enabled.
# -----
# supportKrb5 no
# -----
 serverKrbPrinc <kerberos identity>
# Default Value: none
 Example:
   serverKrbPrinc LDAP/hostname@REALM
# NOTES:
  If Kerberos support will be enabled then specify the kerberos
   identity that is associated with the LDAP server.
# serverKrbPrinc LDAP/host.endicott.ibm.com@REALM
 krbKeytab <filename>
# Default Value: none
# Example:
   krbKeytab /etc/ldap/ldapkrbkeytab
   krbKeytab LDAPKRBKEYTAB
# NOTES:
  If the Kerberos KDC does not exist on the same machine as the LDAP
  server then you must specify the key table file you created for the
   server. The value entered will then be set in the KRB5 KTNAME
   environment variable. However if the KDC and LDAP server reside
   on the same machine a key table is not necessary. Simply specify
   none and the KRB5 SERVER KEYTAB environment variable will be set.
# krbKeytab none
 _____
 krbLDAPAdmin <kerberos identity>
# Default Value: none
# Example:
   krbLDAPAdmin ibm-kn=principal@REALM
   krbLDAPAdmin ibm-kn=adminprincipal@YOURREALM
# NOTES:
   This option allows the LDAP administrator to bind to the server
   via Kerberos and maintain administrative authority.
```

krbLDAPAdmin ibm-kn=adminprincipal@YOURREALM

```
Database-specific SECTIONS
# -----
 SDBM-specific CONFIGURATION SETTINGS
# -----
# database <sdbm|tdbm|rdbm|hcd|exop> <dllname>
# Default Value: none
# Example:
 database sdbm GLDBSDBM
# database sdbm GLDBSDBM
# suffix <toplevelname>
# Default Value: none
# Example:
 suffix "sysplex=sysplex1"
# NOTES:
 This option is REQUIRED when using the SDBM database.
# suffix "sysplex=sysplex1"
# -----
# extendedgroupsearching <on off>
# Default Value: off
# Example:
 extendedgroupsearching on
 Set this option to on if you know that this database contains group
```

```
membership information for users in other databases that are necessary
   for complete authorization checking of a client request.
   By default, this option is set to off. This is useful when your
   LDAP server will be used only for authentication (LDAP BIND) checks.
   When set to off, this option is also useful if you know that this database
   does not contain group membership definitions that will be needed when
   accessing information in other databases, such as TDBM or RDBM.
# krbIdentityMap <on|off>
# Default Value: off
# Example:
   krbIdentityMap on
# NOTES:
   Set this option to on if you want Kerberos identities to be mapped to
   backend entries for access control.
# krbIdentityMap off
# sizeLimit
               <number>
# Default Value: 500
# Example:
  sizeLimit 350
# timeLimit <seconds>
# Default Value: 3600
# Example:
   timeLimit
                60
  TDBM-specific CONFIGURATION SETTINGS
# database <sdbm|tdbm|rdbm|hcd|exop> <dllname>
# Default Value: none
```

```
# Example:
  database tdbm GLDBTDBM
# database tdbm GLDBTDBM
 -----
# suffix <toplevelname>
# Default Value: none
# Example:
  suffix "o=Your Company"
# NOTES:
#
  This option is REQUIRED when using the TDBM database.
# suffix "o=Your Company"
# dsnaoini <filename>
# Default Value: none
# Example:
  dsnaoini /etc/ldap/db2cli.ini
# NOTES:
  The format for this configuration value matches the format required
  by the DB2 Call Level Interface initialization file name. This is
  documented in: DB2 for OS/390 V5 Call Level Interface Guide and Reference
  ( SC26-8959-02). The DB2 initialization file can also be specified
   using the DSNAOINI environment variable or by defining a DSNAOINI
   DD card in the JCL used to start the LDAP server.
# dsnaoini /etc/ldap/db2cli.ini
# -----
# servername <name>
# Default Value: none
# Example:
  servername LOC1
# NOTES:
  The value for this configuration option should be the name of one of the
  data sources defined in the DB2 initialization file (see the dsnaoini
  configuration option). This name is also referred to as the location name
   in some DB2 documentation.
  This option is REQUIRED when using the TDBM database.
 -----
# servername LOC1
# dbuserid <userid>
# Default Value: none
```

```
# Example:
   dbuserid LDAPSRV
   This configuration option is used to indicate the names of the DB2 tables
   that are used by the TDBM database. DB2 table names are prefixed with the
   userid that was used to create the tables.
   The dbuserid option must specify the userid under which the DB2
   tables that the LDAP server uses were created. If ldif2tdbm is
   used to prime the LDAP Directory, then the userid specified
   here should be the userid that ran the ldif2tdbm command.
   This option is REQUIRED when using the TDBM database.
# -----
# dbuserid LDAPSRV
 databasename <name>
# Default Value: none
# Example:
   databasename LDAPR10
   The databasename option must match the name of the database
   that was created to hold the tables that the LDAP server uses.
   This option is REQUIRED when using the TDBM database.
# databasename LDAPR10
 attroverflowsize <number>
# Default Value: 255
# Example:
   attroverflowsize 10000
   The attroverflowsize option indicates the maximum length of an
   entry attribute value that will be stored with the rest of the
   entry information. Accessing attribute values that exceed this
   size will incur a performance penalty (extra table accesses
   during retrieval). Attribute values that have lengths below
   this value will be stored with the rest of the entry information
   and will be retrieved when the entry is retrieved. Thus,
   setting this value to a low value may cause un-requested data
   to be retrieved unnecessarily.
 -----
# pwEncryption <none|crypt|md5|sha|des:keylabel>
# Default Value: none
# Example:
```

```
pwEncryption crypt
 extendedgroupsearching <on off>
# Default Value: off
# Example:
   extendedgroupsearching on
# NOTES:
  Set this option to on if you know that this database contains group
   membership information for users in other databases that are necessary
   for complete authorization checking of a client request.
#
   By default, this option is set to off. This is useful when your
   LDAP server will be used only for authentication (LDAP BIND) checks.
   When set to off, this option is also useful if you know that this database
   does not contain group membership definitions that will be needed when
   accessing information in other databases, such as TDBM or RDBM.
# krbIdentityMap <on|off>
# Default Value: off
# Example:
  krbIdentityMap on
# NOTES:
   Set this option to on if you want Kerberos identities to be mapped to
   backend entries for access control.
# krbIdentityMap off
# ______
# useNativeAuth <selected all off>
# Default Value: off
# Example:
  useNativeAuth all
# NOTES:
   This option enables native authentication in the TDBM
   backend. If the value is:
       SELECTED - only entries within native subtrees with the ibm-nativeId
                  attribute will use native authentication.
       ALL
                - all entries within native subtrees will use native
                  authentication. These entries can contain the ibm-nativeId
                 or uid attribute to specify the RACF ID.
                - no entries will participate in native authentication
# useNativeAuth off
# nativeUpdateAllowed <on|yes|off|no>
```

```
# Default Value: no
 Example:
   nativeUpdateAllowed yes
   This option enables native password changes in RACF to occur
   through the TDBM backend if the useNativeAuth option is specified.
# nativeUpdateAllowed no
# nativeAuthSubtree <all|distinguishedname>
# Example:
   nativeAuthSubtree ou=pok,o=IBM,c=US
   nativeAuthSubtree ou=endicott,o=IBM,c=US
# NOTES:
  This option specifies the distinguished name of a subtree
   where all of its entries participate in Native Authentication.
   This paramter can appear zero or more time to specify all
   subtrees that will use native authentication.
   If this parameter is omitted, contains no value, or is set
   to 'all' then the entire directory will be subject to
   native authentication. This option is only valid if the
   useNativeAuth option is also specified.
 -----
# sizeLimit <number>
# Default Value: 500
# Example:
  sizeLimit 350
 _____
# -----
# timeLimit <seconds>
# Default Value: 3600
# Example:
   timeLimit
               60
# ______
# readonly <on|off>
# Default Value: off
# Example:
   readonly on
# NOTES:
   Use this option to turn off all updates (adds, modifies, or deletes)
   for a database. Do NOT use this option in replica servers as this will
```

```
disallow replica updates sent from the master server.
# -----
# masterserver <1dapURL>
# Default Value: none
# Example:
  masterserver ldap://ldap.yourcompany.com
# NOTES:
  Use this option on a replica server to refer clients to the master
  server when update requests are made to the replica server. This option
  is different from the referral option which is used to point clients
  to different portions of the directory name-space.
# -----
# masterserverDN <distinguishedname>
# Default Value: none
# Example:
   masterserverDN "cn=Master Server, o=Your Company"
# NOTES:
  Use this option on a replica server to define the distinguished name
  that the master replica server will use to authenticate itself to the
  replica server in order to forward updates to the replica server.
# masterserverPW <password>
# Default Value: none
# Example:
  masterserverPW secret
  This configuration option is optional and NOT recommended for general
  use. It is recommended that this option be used only during the initial
  startup of the LDAP server only until an entry in the LDAP directory
  can be defined which will serve as the master server identity.
 -----
# multiserver y
# Default Value: n
# Example:
  multiserver y
# NOTES:
   This configuration option must be specified as 'y' if multiple LDAP
   servers will be accessing the same DB2 tables simultaneously.
```

```
RDBM-specific CONFIGURATION SETTINGS
 -----
 database <sdbm|tdbm|rdbm|hcd|exop> <dllname>
# Default Value: none
# Example:
   database rdbm GLDBRDBM
# database rdbm GLDBRDBM
# suffix <toplevelname>
# Default Value: none
# Example:
   suffix "o=Your Other Company"
# NOTES:
   This option is REQUIRED when using the TDBM database.
   More than one value can be provided. "cn=localhost" MUST be specified
   as one of the suffixes for the RDBM database. No suffix value can
   overlap with a suffix setting in another database.
   This option is REQUIRED when using the RDBM database.
# suffix "cn=localhost"
# suffix "o=Your Other Company"
# dsnaoini <filename>
# Default Value: none
 Example:
   dsnaoini /etc/ldap/db2cli.ini
   The format for this configuration value matches the format required
   by the DB2 Call Level Interface initialization file name. This is
   documented in: DB2 for OS/390 V5 Call Level Interface Guide and Reference
   ( SC26-8959-02). The DB2 initialization file can also be specified
   using the DSNAOINI environment variable or by defining a DSNAOINI
   DD card in the JCL used to start the LDAP server.
   This option is REQUIRED when using the RDBM database. If both the TDBM
```

```
and RDBM databases will be used simultaneously, this option must be set
   to the same value in both database definitions.
# dsnaoini /etc/ldap/db2cli.ini
 -----
# servername <name>
# Default Value: none
# Example:
  servername LOC1
# NOTES:
  The value for this configuration option should be the name of one of the
  data sources defined in the DB2 initialization file (see the dsnaoini
   configuration option). This name is also referred to as the location name
   in some DB2 documentation.
   This option is REQUIRED when using the RDBM database.
# servername LOC1
# dbuserid <userid>
# Default Value: none
# Example:
  dbuserid USER1
# NOTES:
   This configuration option is used to indicate the names of the DB2 tables that are used by the TDBM database. DB2 table names are prefixed with the
   userid that was used to create the tables.
   This option is REQUIRED when using the RDBM database.
# -----
# dbuserid USER1
# databasename <name>
# Default Value: none
# Example:
  databasename LDAPRDBM
# NOTES:
   The databasename option must match the name of the database
   that was created to hold the tables that the LDAP server uses.
   This option is REQUIRED when using the RDBM database.
# databasename LDAPRDBM
 -----
# tbspaceentry <name>
```

```
# Default Value: none
# Example:
   tbspaceentry ldapent
# NOTES:
   The name value for this configuration option must match the
   tablespace name of the first 4k tablespace that is defined in the SPUFI
   script that is run to define the DB2 tablespaces for the RDBM
   database to use.
   This option is REQUIRED when using the RDBM database.
# tbspaceentry ldapent
# -----
 tbspace32k <name>
# Default Value: none
# Example:
   tbspace32k ldap32k
# NOTES:
   The name value for this configuration option must match the
   tablespace name of the 32k tablespace that is defined in the SPUFI
   script that is run to define the DB2 tablespaces for the RDBM
   database to use.
  This option is REQUIRED when using the RDBM database.
# tbspace32k ldap32k
# tbspace4k <name>
# Default Value: none
# Example:
   tbspace4k ldap4k
# NOTES:
   The name value for this configuration option must match the
   tablespace name of the second 4k tablespace that is defined in the SPUFI
   script that is run to define the DB2 tablespaces for the RDBM
   database to use.
   This option is REQUIRED when using the RDBM database.
# tbspace4k ldap4k
# -----
# tbspacemutex <name>
# Default Value: none
# Example:
   tbspacemutex ldapmtx
  The name value for this configuration option must match the
```

```
tablespace name of the third 4k tablespace that is defined in the SPUFI
   script that is run to define the DB2 tablespaces for the RDBM
   database to use.
  This option is REQUIRED when using the RDBM database.
# tbspacemutex ldapmtx
# extendedgroupsearching <on off>
# Default Value: off
# Example:
   extendedgroupsearching on
# NOTES:
   Set this option to on if you know that this database contains group
   membership information for users in other databases that are necessary
   for complete authorization checking of a client request.
  By default, this option is set to off. This is useful when your
   LDAP server will be used only for authentication (LDAP BIND) checks.
   When set to off, this option is also useful if you know that this database
   does not contain group membership definitions that will be needed when
   accessing information in other databases, such as TDBM or RDBM.
 sizeLimit
             <number>
# Default Value: 500
# Example:
  sizeLimit
             350
# -----
# timeLimit <seconds>
# Default Value: 3600
# Example:
  timeLimit
               60
 _____
 readonly <on|off>
# Default Value: off
# Example:
   readonly on
# NOTES:
  Use this option to turn off all updates (adds, modifies, or deletes)
   for a database. Do NOT use this option in replica servers as this will
   disallow replica updates sent from the master server.
```

```
# masterserver <1dapURL>
# Default Value: none
# Example:
   masterserver ldap://ldap.yourcompany.com
# NOTES:
   Use this option on a replica server to refer clients to the master
   server when update requests are made to the replica server. This option
   is different from the referral option which is used to point clients
   to different portions of the directory name-space.
# masterserverDN <distinguishedname>
# Default Value: none
# Example:
   masterserverDN "cn=Master Server, o=Your Company"
   Use this option on a replica server to define the distinguished name
   that the master replica server will use to authenticate itself to the
   replica server in order to forward updates to the replica server.
# masterserverPW <password>
# Default Value: none
 Example:
   masterserverPW secret
# NOTES:
   This configuration option is optional and NOT recommended for general
   use. It is recommended that this option be used only during the initial
   startup of the LDAP server only until an entry in the LDAP directory
   can be defined which will serve as the master server identity.
 -----
# multiserver y
# Default Value: n
 Example:
   multiserver y
# NOTES:
   This configuration option must be specified as 'y' if multiple LDAP
   servers will be accessing the same DB2 tables simultaneously.
```

```
-----
# pwEncryption <none|crypt|md5|sha|des:keylabel>
# Default Value: none
# Example:
   pwEncryption crypt
 _____
# index <default|attname> <eq[,approx[,sub]]|none>
# Default Value: none
# Example:
 index default eq
index cn eq
index ou eq
index sn eq
index userCertificate none
#
# NOTES:
  This configuration option can be specified multiple times in the
   RDBM database section.
# index cn eq
# index sn eq
 EXOP-specific CONFIGURATION SETTINGS
# database <sdbm|tdbm|rdbm|hcd|exop> <dllname>
# Default Value: none
# Example:
  database exop GLDXPDIR
# database exop GLDXPDIR
END OF CONFIGURATION FILE
```

Appendix B. Minimum schema for TDBM

This appendix shows the minimum schema for TDBM assuming the suffix is o=ibm us, c=us.

```
cn=schema, o=ibm us, c=us
cn=SCHEMA
subtreespecification=NULL
objectclass=TOP
objectclass=SUBSCHEMA
objectclass=SUBENTRY
objectclass=IBMSUBSCHEMA
attributetypes=( 1.3.18.0.2.4.285 NAME 'aclentry' EQUALITY caseexactmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{32700} USAGE directoryoperation )
attributetypes=( 1.3.18.0.2.4.286 NAME 'aclpropagate' SYNTAX 1.3.6.1.4.1.1466.115.121.1.7{5} USAGE directoryoperation SINGLE-VALUE )
attributetypes=( 1.3.18.0.2.4.287 NAME 'aclsource'
SYNTAX 1.3.6.1.4.1.1466.115.121.1.12{1000} USAGE directoryoperation SINGLE-VALUE NO-USER-MODIFICATION )
attributetypes=( 1.3.6.1.4.1.1466.101.120.6 NAME 'altserver SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE dsaoperation )
attribute types = (\ 2.5.21.5\ NAME\ 'attribute types'\ EQUALITY\ object identifier first component match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.3 USAGE directoryoperation )
attributetypes=( 2.5.4.3 NAME ( 'cn' 'commonname' ) SUP name ) attributetypes=( 2.5.18.1 NAME 'createtimestamp' EQUALITY generalizedtimematch ORDERING generalizedtimeorderingmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24 SINGLE-VALUE NO-USER-MODIFICATION USAGE directoryoperation )
attributetypes=( 2.5.18.3 NAME 'creatorsname' EQUALITY distinguishednamematch SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 SINGLE-VALUE NO-USER-MODIFICATION USAGE directoryoperation )
attributetypes=( 2.5.21.2 NAME 'ditcontentrules' EQUALITY objectidentifierfirstcomponentmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.16 USAGE directoryoperation )
attributetypes=( 2.5.21.1 NAME 'ditstructurerules' EQUALITY integerfirstcomponentmatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.17 USAGE directoryoperation ) attributetypes=( 2.5.4.49 NAME ( 'dn' 'distinguishedname' ) EQUALITY distinguishednamematch SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 )
attributetypes=( 1.3.18.0.2.4.288 NAME 'entryowner' EQUALITY caseexactmatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{1000} USAGE directoryoperation ) attributetypes=( 1.3.18.0.2.4.470 NAME 'ibmattributetypes' EQUALITY objectidentifierfirstcomponentmatch
SYNTAX 1.3.18.0.2.8.1 USAGE directoryoperation ) attributetypes=( 1.3.18.0.2.4.1780 NAME 'ibm-entryuuid' EQUALITY ibm-entryuuidmatch
  SYNTAX 1.3.18.0.2.8.1 USAGE dsaoperation NO-USER-MODIFICATION SINGLE-VALUE DESC
'Uniquely identifies an ldap entry throughout its life.' ) attributetypes=( 0.9.2342.19200300.100.1.24 NAME 'lastmodifiedby'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12{1000} SINGLE-VALUE )
attributetypes=( 0.9.2342.19200300.100.1.23 NAME 'lastmodifiedtime'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24{30} SINGLE-VALUE )
attributetypes=( 1.3.6.1.4.1.1466.101.120.16 NAME 'ldapsyntaxes' EQUALITY objectidentifierfirstcomponentmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.54 USAGE directoryoperation )
attributetypes=( 2.5.21.4 NAME 'matchingrules' EQUALITY objectidentifierfirstcomponentmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.30 USAGE directoryoperation )
attributetypes=( 2.5.21.8 NAME 'matchingruleuse' EQUALITY objectidentifierfirstcomponentmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.31 USAGE directoryoperation )
attributetypes=( 2.5.4.31 NAME 'member' SUP distinguishedname ) attributetypes=( 2.5.18.4 NAME 'modifiersname' EQUALITY distinguishednamematch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 SINGLE-VALUE NO-USER-MODIFICATION USAGE directoryoperation )
attributetypes=( 2.5.18.2 NAME 'modifytimestamp' EQUALITY generalizedtimematch ORDERING generalizedtimeorderingmatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.24 SINGLE-VALUE NO-USER-MODIFICATION USAGE directoryoperation ) attributetypes=( 2.5.4.41 NAME 'name' EQUALITY caseignorematch SUBSTR caseignoresubstringsmatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{32768} ) attributetypes=( 2.5.21.7 NAME 'nameforms' EQUALITY objectidentifierfirstcomponentmatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.35 USAGE directoryoperation )
attributetypes=( 1.3.6.1.4.1.1466.101.120.5 NAME 'namingcontexts
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 USAGE dsaoperation )
attributetypes=( 2.5.4.0 NAME 'objectclass' EQUALITY objectidentifiermatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 )
attributetypes=( 2.5.21.6 NAME 'objectclasses' EQUALITY objectidentifierfirstcomponentmatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.37 USAGE directoryoperation ) attributetypes=( 1.3.18.0.2.4.289 NAME 'ownerpropagate'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.7{5} USAGE directoryoperation SINGLE-VALUE )
attributetypes=( 1.3.18.0.2.4.290 NAME 'ownersource
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12{1000} USAGE directoryoperation SINGLE-VALUE NO-USER-MODIFICATION )
attributetypes=( 2.16.840.1.113730.3.1.34 NAME 'ref
attributetypes=( 2.10.040.1.113/30.3.1.34 NAME | 161

SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{100} EQUALITY caseexactmatch )

attributetypes=( 1.3.18.0.2.4.299 NAME 'replicabinddn'

SYNTAX 1.3.6.1.4.1.1466.115.121.1.12{1000} USAGE directoryoperation )
attributetypes=( 1.3.18.0.2.4.302 NAME 'replicabindmethod' SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{100} USAGE directoryoperation ) attributetypes=( 1.3.18.0.2.4.300 NAME ( 'replicacredentials' 'replicabindcredentials' )
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.5{128} SINGLE-VALUE USAGE directoryoperation )
attributetypes=( 1.3.18.0.2.4.298 NAME 'replicahost'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{100} SINGLE-VALUE USAGE directoryoperation )
attributetypes=( 1.3.18.0.2.4.301 NAME 'replicaport'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{10} SINGLE-VALUE USAGE directoryoperation )
```

Minimum Schema for TDBM

```
attributetypes=( 1.3.18.0.2.4.304 NAME 'replicaupdatetimeinterval'
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15\{20\} SINGLE-VALUE USAGE directoryoperation ) attributetypes=( 1.3.18.0.2.4.303 NAME 'replicausess1'
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{10} SINGLE-VALUE USAGE directoryoperation )
attributetypes=( 2.5.18.10 NAME 'subschemasubentry' EQUALITY distinguishednamematch SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 SINGLE-VALUE NO-USER-MODIFICATION USAGE directoryoperation )
attributetypes=( 2.5.18.6 NAME 'subtreespecification
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE USAGE directoryOperation )
attributetypes=( 1.3.6.1.4.1.1466.101.120.13 NAME 'supportedcontrol
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 USAGE dsaoperation )
attributetypes=( 1.3.6.1.4.1.1466.101.120.7 NAME 'supportedextension'
SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 USAGE dsaoperation ) attributetypes=( 1.3.6.1.4.1.1466.101.120.15 NAME 'supported!dapversion' SYNTAX 1.3.6.1.4.1.1466.115.121.1.27 USAGE dsaoperation )
attributetypes=( 1.3.6.1.4.1.1466.101.120.14 NAME 'supportedsaslmechanisms'
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 USAGE dsaoperation )
attributetypes=( 2.5.4.35 NAME 'userpassword' EQUALITY octetstringmatch
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.40{128} )
ibmattributetypes=( 1.3.18.0.2.4.285 ACCESS-CLASS restricted
ibmattributetypes=( 1.3.18.0.2.4.286 ACCESS-CLASS restricted )
ibmattributetypes=( 1.3.18.0.2.4.287 ACCESS-CLASS system )
ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.6 ACCESS-CLASS normal )
ibmattributetypes=( 2.5.21.5 ACCESS-CLASS system )
ibmattributetypes=( 2.5.4.3 ACCESS-CLASS normal ) ibmattributetypes=( 2.5.18.1 ACCESS-CLASS system )
ibmattributetypes=( 2.5.18.3 ACCESS-CLASS system ) ibmattributetypes=( 2.5.21.2 ACCESS-CLASS system )
ibmattributetypes=( 2.5.21.1 ACCESS-CLASS system ) ibmattributetypes=( 2.5.4.49 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.288 ACCESS-CLASS restricted )
ibmattributetypes=( 1.3.18.0.2.4.470 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.1780 ACCESS-CLASS normal )
ibmattributetypes=( 0.9.2342.19200300.100.1.24 ACCESS-CLASS system )
ibmattributetypes = ( \ 0.9.2342.19200300.100.1.23 \ ACCESS-CLASS \ system \ ) \\ ibmattributetypes = ( \ 1.3.6.1.4.1.1466.101.120.16 \ ACCESS-CLASS \ system \ ) \\
ibmattributetypes=( 2.5.21.4 ACCESS-CLASS system )
ibmattributetypes=( 2.5.21.8 ACCESS-CLASS system )
ibmattributetypes=( 2.5.4.31 ACCESS-CLASS normal ibmattributetypes=( 2.5.18.4 ACCESS-CLASS system
ibmattributetypes=( 2.5.18.2 ACCESS-CLASS system ibmattributetypes=( 2.5.4.41 ACCESS-CLASS normal
ibmattributetypes=( 2.5.21.7 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.5 ACCESS-CLASS normal )
ibmattributetypes=( 2.5.4.0 ACCESS-CLASS normal )
ibmattributetypes=( 2.5.21.6 ACCESS-CLASS system )
ibmattributetypes=( 1.3.18.0.2.4.289 ACCESS-CLASS restricted )
bmattributetypes=( 1.3.18.0.2.4.289 ACCESS-CLASS restricted )
ibmattributetypes=( 1.3.18.0.2.4.290 ACCESS-CLASS system )
ibmattributetypes=( 2.16.840.1.113730.3.1.34 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.299 ACCESS-CLASS critical )
ibmattributetypes=( 1.3.18.0.2.4.302 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.300 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.298 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.298 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.301 ACCESS-CLASS normal
ibmattributetypes=( 1.3.18.0.2.4.304 ACCESS-CLASS normal
ibmattributetypes=( 1.3.18.0.2.4.303 ACCESS-CLASS normal )
ibmattributetypes=( 2.5.18.10 ACCESS-CLASS system )
ibmattributetypes=( 2.5.18.6 ACCESS-CLASS system )
ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.13 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.7 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.15 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.14 ACCESS-CLASS normal )
ibmattributetypes=( 2.5.4.35 ACCESS-CLASS critical )
objectclasses=( 1.3.18.0.2.6.174 NAME 'ibmsubschema' SUP top AUXILIARY MAY ibmattributetypes )
objectclasses=( 2.16.840.1.113730.3.2.6 NAME 'referral' SUP top MUST ref ) objectclasses=( 2.5.20.1 NAME 'subschema' SUP top
   AUXILIARY MAY ( ditstructurerules $ nameforms $ ditcontentrules $ objectclasses $
   attributetypes $ matchingrules $ matchingruleuse $ ldapsyntaxes ) )
objectclasses=( 2.5.17.0\, NAME 'subentry' SUP top STRUCTURAL MUST ( cn \, subtreespecification ) ) objectclasses=( 2.5.6.0\, NAME 'top' ABSTRACT MUST objectclass )
| dapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.3 DESC | dapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.5 DESC
                                                                       'attribute type description' )
                                                                        'binary' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.7 DESC ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.15 DESC
                                                                       'boolean' )
                                                                         'directory string' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.12 DESC
                                                                         'dn' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.16 DESC
                                                                         'dit content rule description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.17 DESC
                                                                         'dit structure rule description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.24 DESC 'generalized time') ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.26 DESC 'ia5 string')
ldapsyntaxes=( 1.3.18.0.2.8.1 DESC 'ibm attribute type description' )
ldapsyntaxes=( 1.3.18.0.2.8.3 DESC '16 byte DCE UUID value in hex string representation.
    TimeLow-TimeMid-Version/TimeHi-Variant/ClkSegHi/ClkSegLow-NodeID.
   For example, 11fdfa9e-1ec6-11d5-8ace-0006292197c7'
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.27 DESC 'integer' )
```

Minimum Schema for TDBM

```
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.54 DESC
                                                      'ldap syntax description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.30 DESC
                                                      'matching rule description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.31 DESC
                                                      'matching rule use description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.35 DESC
                                                      'name form description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.37 DESC
                                                      'object class description' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.40 DESC
                                                      'octet string' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.38 DESC
                                                      'oid' )
                                                      'substring assertion' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.58 DESC
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.50 DESC ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.53 DESC
                                                      'telephone number' )
                                                      'utc time' )
matchingrules=( 2.5.13.13 NAME 'booleanmatch'
 SYNTAX 1.3.6.1.4.1.1466.115.121.1.7 )
matchingrules=( 1.3.6.1.4.1.1466.109.114.1 NAME 'caseexactia5match'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 )
matchingrules=( 2.5.13.5 NAME 'caseexactmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )
matchingrules=( 2.5.13.6 NAME 'caseexactorderingmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )
matchingrules=( 2.5.13.7 NAME 'caseexactsubstringsmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.58 )
matchingrules=( 1.3.6.1.4.1.1466.109.114.2 NAME 'caseignoreia5match'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 )
matchingrules=( 2.5.13.3 NAME 'caseignoreorderingmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )
matchingrules=( 2.5.13.2 NAME 'caseignorematch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )
matchingrules=( 2.5.13.4 NAME 'caseignoresubstringsmatch'
SYNTAX 1.3.6.1.4.1.1466.115.121.1.58 ) matchingrules=( 2.5.13.1 NAME 'distinguishednamematch'
SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 ) matchingrules=( 1.3.18.0.2.4.405 NAME 'distinguishednameorderingmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.12 )
matchingrules=( 1.3.18.0.2.22.2 NAME 'ibm-entryuuidmatch'
  SYNTAX 1.3.18.0.2.8.3 )
matchingrules=( 2.5.13.27 NAME 'generalizedtimematch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24 )
matchingrules=( 2.5.13.28 NAME 'generalizedtimeorderingmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24 )
matchingrules=( 2.5.13.14 NAME 'integermatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27 )
matchingrules=( 2.5.13.29 NAME 'integerfirstcomponentmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27 )
matchingrules=( 2.5.13.0 NAME 'objectidentifiermatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 )
matchingrules=( 2.5.13.30 NAME 'objectidentifierfirstcomponentmatch'
 SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 )
matchingrules=( 2.5.13.17 NAME 'octetstringmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.40 )
matchingrules=( 2.5.13.20 NAME 'telephonenumbermatch'
 SYNTAX 1.3.6.1.4.1.1466.115.121.1.50 )
matchingrules=( 2.5.13.21 NAME 'telephonenumbersubstringsmatch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.58 )
matchingrules=( 2.5.13.25 NAME 'utctimematch'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.53 )
```

Appendix C. SDBM schema

This appendix shows the schema for SDBM assuming the suffix is sysplex=myRACF.

```
CN=SCHEMA, sysplex=myRACF
cn=SCHEMA
subtreespecification=NULL
objectclass=TOP
objectclass=SUBSCHEMA
objectclass=SUBENTRY
objectclass=IBMSUBSCHEMA
 attributetypes=( 2.5.21.5 NAME 'attributeTypes' EQUALITY objectIdentifierFirstComponentMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.3 USAGE directoryOperation ) attributetypes=( 1.3.18.0.2.4.470 NAME 'ibmAttributeTypes' EQUALITY objectIdentifierFirstComponentMatch
SYNTAX 1.3.18.0.2.8.1 USAGE directoryOperation )
attributetypes=( 1.3.6.1.4.1.1466.101.120.16 NAME 'ldapSyntaxes' EQUALITY objectIdentifierFirstComponentMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.54 USAGE directoryOperation )
attributetypes=( 2.5.21.4 NAME 'matchingRules' EQUALITY objectIdentifierFirstComponentMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.30 USAGE directoryOperation )
SYNIAX 1.3.6.1.4.1.1466.115.121.1.30 USAGE directoryOperation )
attributetypes=( 2.5.21.8 NAME 'matchingRuleUse' EQUALITY objectIdentifierFirstComponentMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.31 USAGE directoryOperation )
attributetypes=( 2.5.4.0 NAME 'objectClass' EQUALITY objectIdentifiermatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 )
attributetypes=( 2.5.21.6 NAME 'objectClasses' EQUALITY objectIdentifierFirstComponentMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.37 USAGE directoryOperation )
attributetypes=( 2.5.21.6 NAME 'objectClasses' Equality objectIdentifierFirstComponentMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.37 USAGE directoryOperation )
attributetypes=( 2.5.18.6 NAME 'subtreeSpecification'
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE USAGE directoryOperation )
attributetypes=( 2.5.21.2 NAME 'ditContentRules' EQUALITY objectIdentifierFirstComponentMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.16 USAGE directoryOperation )
attribute types = (\ 2.5.21.1\ NAME\ 'ditStructure Rules'\ EQUALITY\ integer First Component Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.17 USAGE directoryOperation ) attributetypes=( 2.5.21.7 NAME 'nameForms' EQUALITY objectIdentifierFirstComponentMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.35 USAGE directoryOperation )
attributetypes= ( 2.5.4.41 NAME 'name' EQUALITY caseIgnoreMatch SUBSTR caseIgnoreSubstringsMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.15{32768} )
attributetypes=(2.5.4.3 NAME ('cn' 'commonname') SUP name) attributetypes=(1.3.18.0.2.4.185 NAME 'sysplex' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications)
attributetypes=( 1.3.18.0.2.4.186 NAME 'profileType' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.188 NAME 'racfAuthorizationDate' DESC 'Date is displayed in yy.ddd format.' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.189 NAME 'racfOwner' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.190 NAME 'racfInstallationData' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.191 NAME 'racfDatasetModel' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.187 NAME 'racfid' DESC 'Identifies the name of a 0S/390 Security Server userid or groupid.'
EQUALITY caseIgnore2 IASMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.197 NAME 'racfAttributes' EQUALITY caseIgnoreIASMatch
     SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.198 NAME 'racfPassword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.200 NAME 'racfPasswordChangeDate' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.199 NAME 'racfPasswordInterval' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.201 NAME 'racfProgrammerName' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.202 NAME 'racfDefaultGroup' EQUALITY caseIgnoreIA5Match
attributetypes=( 1.3.18.0.2.4.202 NAME 'racfletaultGroup' EQUALITY caseIgnoreIA5Matcn SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications) attributetypes=( 1.3.18.0.2.4.203 NAME 'racflastAccess' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications) attributetypes=( 1.3.18.0.2.4.214 NAME 'racfSecurityLabel' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications) attributetypes=( 1.3.18.0.2.4.205 NAME 'racfSecurityCategoryList' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications)
     SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
 attributetypes=( 1.3.18.0.2.4.206 NAME 'racfRevokeDate' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.207 NAME 'racfResumeDate' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.208 NAME 'racfLogonDays' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.209 NAME 'racfLogonTime' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.210 NAME 'racfClassName' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.211 NAME 'racfConnectGroupName' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.212 NAME 'racfConnectGroupAuthority' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.213 NAME 'racfConnectGroupUACC' EQUALITY caseIgnoreIA5Match
     SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.204 NAME 'racfSecurityLevel' EQUALITY caseIgnoreIA5Match
```

```
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attribute types = ( \ 1.3.18.0.2.4.192 \ NAME \ 'racf Superior Group' \ EQUALITY \ case Ignore IA5 Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.193 NAME 'racfGroupNoTermUAC' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.194 NAME 'racfSubGroupName' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.196 NAME 'racfGroupUserAccess' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.195 NAME 'racfGroupUserids' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.215 NAME 'SAFDfpDataApplication' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
                                              'SAFDfpDataClass' EQUALITY caseIgnoreIA5Match
attributetypes=( 1.3.18.0.2.4.216 NAME
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.217 NAME 'SAFDfpManagementClass' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attribute types = (\ 1.3.18.0.2.4.218\ NAME\ 'SAFD fpStorageClass'\ EQUALITY\ caseIgnore IA5 Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.219 NAME 'racfOmvsGroupId' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.220 NAME 'racf0vmGroupId' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.265 NAME 'racfOmvsUid' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.266 NAME 'racfOmvsHome' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.267 NAME 'racfOmvsInitialProgram' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.280 NAME 'racfOvmUid' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.281 NAME 'racfOvmHome' EQUALITY caseExactMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.282 NAME 'racf0vmInitialProgram' EQUALITY caseExactMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attribute types = ( \ 1.3.18.0.2.4.283 \ NAME \ 'racfOvmFileSystemRoot' \ EQUALITY \ caseExactMatch') \\
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.221 NAME 'SAFAccountNumber' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.223 NAME
                                              'SAFDestination' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.224 NAME 'SAFHoldClass' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) ttributetypes=( 1.3.18.0.2.4.225 NAME 'SAFJobClass' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.225 NAME
                                              'SAFMessageClass' EQUALITY caseIgnoreIA5Match
attributetypes=( 1.3.18.0.2.4.226 NAME
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.227 NAME 'SAFDefaultLoginProc' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.228 NAME 'SAFLogonSize' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.229 NAME
                                              'SAFMaximumRegionSize' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.230 NAME 'SAFDefaultSysoutClass' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.231 NAME
                                               'SAFUserData' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.232 NAME
                                              'SAFDefaultUnit' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.233 NAME 'SAFTsoSecurityLabel' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.222 NAME 'SAFDefaultCommand' EQUALITY caseExactMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.237 NAME 'racfOperatorClass' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.236 NAME 'racfOperatorIdentification' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.238 NAME 'racfOperatorPriority' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.239 NAME 'racfOperatorReSignon' EOUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.240 NAME 'racfTerminalTimeout' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
                                              'racfStorageKeyword' EQUALITY caseIgnoreIA5Match
attributetypes=( 1.3.18.0.2.4.241 NAME
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.242 NAME 'racfAuthKeyword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.243 NAME 'racfMformKeyword' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.244 NAME 'racfLevelKeyword' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.245 NAME 'racfMonitorKeyword' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.246 NAME 'racfRoutcodeKeyword' EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.247 NAME 'racflogCommandResponseKeyword' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.248 NAME 'racfMGIDKeyword' EQUALITY caseIgnoreIA5Match
```

```
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.249 NAME 'racfDOMKeyword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.250 NAME 'racfKEYKeyword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.251 NAME 'racfCMDSYSKeyword' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.252 NAME 'racfUDKeyword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.253 NAME 'racfMscopeSystems' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.254 NAME 'racfAltGroupKeyword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.255 NAME 'racfAutoKeyword' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.234 NAME 'racfPrimaryLanguage' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications
attributetypes=( 1.3.18.0.2.4.235 NAME 'racfSecondaryLanguage' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attribute types \verb== ( 1.3.18.0.2.4.256 \ NAME \ 'racfWork attrUsername' \ EQUALITY \ case \verb== ExactMatch | ExactMatch | Equality | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | ExactMatch | Ex
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.257 NAME 'racfBuilding' EQUALITY caseExactMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.258 NAME 'racfDepartment' EQUALITY caseExactMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.259 NAME 'racfRoom' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.261 NAME 'racfAddressLinel' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.262 NAME 'racfAddressLine2' EQUALITY caseExactMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.263 NAME 'racfAddressLine3' EQUALITY caseExactMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.264 NAME 'racfAddressLine4' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.260 NAME 'racfWorkAttrAccountNumber' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.268 NAME 'racfNetviewInitialCommand' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.269 NAME 'racfDefaultConsoleName' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.270 NAME 'racfCTLKeyword' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.271 NAME 'racfMSGRCVRKeyword' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.272 NAME 'racfNetviewOperatorClass' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications) attributetypes=( 1.3.18.0.2.4.273 NAME 'racfDomains' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.274 NAME 'racfNGMFADMKeyword' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.279 NAME 'racfDCEAutoLogin' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.277 NAME 'racfDCEHomeCell' EQUALITY caseExactMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.278 NAME 'racfDCEHomeCellUUID' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.276 NAME 'racfDCEPrincipal' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.275 NAME 'racfDCEUUID' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.826 NAME 'racfOmvsMaximumAddressSpaceSize' DESC 'Represents the ASSIZEMAX (address-space-size) field of the OMVS RACF SEGMENT.' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.827 NAME 'racfOmvsMaximumCPUTime' DESC 'Represents the CPUTIMEMAX
    (cpu-time) field of the OMVS RACF SEGMENT.' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications
attributetypes-( 1.3.18.0.2.4.828 NAME 'racfOmvsMaximumFilesPerProcess' DESC 'Represents the MMAPAREAMAX (memory-map-size) field of the OMVS RACF SEGMENT.' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.829 NAME 'racfOmvsMaximumMemoryMapArea' DESC 'Represents the ASSIZEMAX
    (address-space-size) field of the OMVS RACF SEGMENT.' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userapplications ) attributetypes=( 1.3.18.0.2.4.830 NAME 'racfOmvsMaximumProcessesPerUID' DESC 'Represents the PROCUSERMAX
    (processes-per-UID) field of the OMVS RACF SEGMENT.' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.831 NAME 'racfOmvsMaximumThreadsPerProcess' DESC 'Represents the THREADSMAX (threads-per-process) field of the OMVS RACF SEGMENT.' EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1099 NAME 'racfLNotesShortName' DESC 'represents the SNAME field of the RACF
LNOTES segment' EQUALITY caseExactMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1100 NAME 'racfNDSUserName' DESC 'Represents the UNAME field of the RACF NDS segment' EQUALITY caseExactMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1153 NAME 'racfCurKeyVersion' DESC 'Current key version' EQUALITY caseIgnoreMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
\label{eq:attributetypes} attributetypes=( 1.2.840.113556.1.4.656 \ \text{NAME 'krbPrincipalName' EQUALITY caseExactMatch SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.2.840.113556.1.4.77 \ \text{NAME 'maxTicketAge'}
```

```
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27 SINGLE-VALUE USAGE userApplications
attributetypes=( 1.3.18.0.2.4.1144 NAME 'racfConnectAttributes' DESC 'RACF Connect Attributes' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 USAGE userApplications) attributetypes=( 1.3.18.0.2.4.1145 NAME 'racfConnectAuthDate' DESC 'RACF Connect Auth Date' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.1146 NAME 'racfConnectCount' DESC 'RACF Connect Count' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1147 NAME 'racfConnectLastConnect' DESC 'RACF Connect Last Connect' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.1148 NAME 'racfConnectOwner' DESC 'RACF Connect Owner' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1149 NAME 'racfConnectResumeDate' DESC 'RACF Connect Resume Date' EQUALITY caseIgnoreIA5Match
SYNTAX 1.3.6.1.4.1.466.115.121.1.26 SINGLE-VALUE USAGE user-Applications )
attributetypes=( 1.3.18.0.2.4.1150 NAME 'racfConnectRevokeDate' DESC 'RACF Connect Revoke Date' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications
attributetypes=( 1.3.18.0.2.4.1151 NAME 'racfGroupId' DESC 'RACF group ID' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1913 NAME 'racfGroupUniversal' DESC 'RACF universal group indicator' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.2007 NAME 'racfEncryptType' DESC 'RACF encrypt type' EQUALITY caseIgnoreIA5Match SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.1152 NAME 'racfUserId' DESC 'RACF userid' EQUALITY caseIgnoreIA5Match
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1162 NAME 'racfLDAPBindDN' DESC 'RACF LDAP Bind DN' EQUALITY caseExactMatch
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.1163 NAME 'racfLDAPBindPw' DESC 'RACF LDAP Bind Password' EQUALITY caseExactMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications ) attributetypes=( 1.3.18.0.2.4.1164 NAME 'racfLDAPHost' DESC 'RACF LDAP Host' EQUALITY caseIgnoreMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.2239 NAME 'racfLDAPProf' DESC 'RACF LDAP Profile Name' EQUALITY caseIgnoreMatch
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.2240 NAME 'racfOmvsGroupIdKeyword' DESC 'RACF group OMVS keyword' EQUALITY caseIgnoreMatch
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
attributetypes=( 1.3.18.0.2.4.2241 NAME 'racfOmvsUidKeyword' DESC 'RACF user OMVS keyword' EQUALITY caseIgnoreMatch
   SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE USAGE userApplications )
ibmattributetypes=( 2.5.21.5 ACCESS-CLASS system ) ibmattributetypes=( 1.3.18.0.2.4.470 ACCESS-CLASS normal ) ibmattributetypes=( 1.3.6.1.4.1.1466.101.120.16 ACCESS-CLASS system ) ibmattributetypes=( 2.5.21.4 ACCESS-CLASS system )
ibmattributetypes=( 2.5.21.8 ACCESS-CLASS system ) ibmattributetypes=( 2.5.4.0 ACCESS-CLASS normal )
ibmattributetypes=( 2.5.21.6 ACCESS-CLASS system ibmattributetypes=( 2.5.18.6 ACCESS-CLASS system
ibmattributetypes=( 2.5.21.2 ACCESS-CLASS system ) ibmattributetypes=( 2.5.21.1 ACCESS-CLASS system ) ibmattributetypes=( 2.5.21.1 ACCESS-CLASS system ) ibmattributetypes=( 2.5.21.7 ACCESS-CLASS normal ) ibmattributetypes=( 2.5.4.41 ACCESS-CLASS normal ) ibmattributetypes=( 2.5.4.3 ACCESS-CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.185 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.186 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.188 ACCESS_CLASS sensitive
ib mattribute types = ( \ 1.3.18.0.2.4.189 \ ACCESS\_CLASS \ sensitive \\ ib mattribute types = ( \ 1.3.18.0.2.4.190 \ ACCESS\_CLASS \ sensitive \\
ibmattributetypes=( 1.3.18.0.2.4.191 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.187 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.197 ACCESS CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.198 ACCESS CLASS critical )
ibmattributetypes=( 1.3.18.0.2.4.200 ACCESS_CLASS critical
ibmattributetypes=( 1.3.18.0.2.4.199 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.201 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.202 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.203 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.214 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.205 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.206 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.207 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.208 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.209 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.210 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.211 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.212 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.213 ACCESS CLASS sensitive ) ibmattributetypes=( 1.3.18.0.2.4.204 ACCESS_CLASS critical )
ibmattributetypes=( 1.3.18.0.2.4.192 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.193 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.194 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.196 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.195 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.215 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.216 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.216 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.217 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.218 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.219 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.220 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.265 ACCESS_CLASS sensitive
ibmattributetypes = (\ 1.3.18.0.2.4.266\ ACCESS\_CLASS\ sensitive\\ ibmattributetypes = (\ 1.3.18.0.2.4.267\ ACCESS\_CLASS\ sensitive
ibmattributetypes=( 1.3.18.0.2.4.280 ACCESS CLASS sensitive )
```

```
ibmattributetypes=( 1.3.18.0.2.4.281 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.282 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.283 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.221 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.223 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.224 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.225 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.226 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.227 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.228 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.229 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.230 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.231 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.232 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.233 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.222 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.237 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.236 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.238 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.239 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.240 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.241 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.242 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.243 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.244 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.245 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.246 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.247 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.248 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.249 ACCESS_CLASS sensitive ibmatt
ibmattributetypes=( 1.3.18.0.2.4.250 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.251 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.252 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.253 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.254 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.255 ACCESS_CLASS sensitive ) ibmattributetypes=( 1.3.18.0.2.4.255 ACCESS_CLASS sensitive ) ibmattributetypes=( 1.3.18.0.2.4.235 ACCESS_CLASS sensitive ) ibmattributetypes=( 1.3.18.0.2.4.256 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.257 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.258 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.259 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.261 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.262 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.263 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.264 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.260 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.268 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.269 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.270 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.271 ACCESS_CLASS sensitive
ibmattributetypes=(\ 1.3.18.0.2.4.272\ ACCESS\_CLASS\ sensitive\\ ibmattributetypes=(\ 1.3.18.0.2.4.273\ ACCESS\_CLASS\ sensitive\\
ibmattributetypes=( 1.3.18.0.2.4.274 ACCESS CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.279 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.277 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.278 ACCESS CLASS sensitive
 ibmattributetypes=( 1.3.18.0.2.4.276 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.275 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.826 ACCESS_CLASS normal )
ibmattributetypes=( 1.3.18.0.2.4.827 ACCESS_CLASS normal ibmattributetypes=( 1.3.18.0.2.4.828 ACCESS_CLASS normal
ibmattributetypes=( 1.3.18.0.2.4.829 ACCESS_CLASS normal
ibmattributetypes=( 1.3.18.0.2.4.830 ACCESS CLASS normal
ibmattributetypes=( 1.3.18.0.2.4.831 ACCESS CLASS normal
ibmattributetypes=( 1.3.18.0.2.4.1099 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1100 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1153 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.2.840.113556.1.4.656 ACCESS_CLASS sensitive )
ibmattributetypes=(\ 1.2.840.113556.1.4.77\ ACCESS\_\overline{C}LASS\ sensitive\ )\\ ibmattributetypes=(\ 1.3.18.0.2.4.1144\ ACCESS\_CLASS\ sensitive\ )
ibmattributetypes=( 1.3.18.0.2.4.1145 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.1146 ACCESS_CLASS sensitive
 ibmattributetypes=( 1.3.18.0.2.4.1147 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1148 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1149 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1150 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.1151 ACCESS_CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1913 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.1913 ACCESS_CLASS sensitive ibmattributetypes=( 1.3.18.0.2.4.2007 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1152 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.1162 ACCESS CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.1163 ACCESS CLASS critical )
ibmattributetypes=( 1.3.18.0.2.4.1164 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.2239 ACCESS CLASS sensitive
ibmattributetypes=( 1.3.18.0.2.4.2240 ACCESS_CLASS sensitive )
ibmattributetypes=( 1.3.18.0.2.4.2241 ACCESS CLASS sensitive )
objectclasses=( 2.5.6.0 NAME 'top' ABSTRACT MUST ( objectclass ) )
```

```
object classes = (\ 2.5.17.0\ NAME\ 'subEntry'\ SUP\ top\ STRUCTURAL\ MUST\ (\ cn\ \$ \ subTreeSpecification\ )\ ) object classes = (\ 2.5.20.1\ NAME\ 'subSchema'\ SUP\ top\ AUXILIARY\ MAY\ (\ ditStructureRules\ \$ \ nameForms\ \$
objectclasses=( 2.5.20.1 NAME 'subSchema' SUP top AUXILIARY MAY ( ditStructureRules $ nameForms $ ditContentRules $ objectclasses $ attributeTypes $ matchingRules $ matchingRuleUse $ ldapSyntaxes ) ) objectclasses=( 1.3.18.0.2.6.174 NAME 'ibmSubSchema' SUP top AUXILIARY MAY ( ibmAttributeTypes ) ) objectclasses=( 1.3.18.0.2.6.55 NAME 'racfbase' DESC 'Represents the base of the Directory Information Tree that publishes information stored by the OS/390 Security Server RACF service.' SUP top STRUCTURAL MAY ( sysplex ) ) objectclasses=( 1.3.18.0.2.6.56 NAME 'racfProfileType' DESC 'Represents a container below which individual RACF profile entries will be published.' SUP top STRUCTURAL MUST ( profileType ) )
objectclasses=( 1.3.18.0.2.6.57 NAME 'racfBaseCommon' DESC 'Represents a common base class for all RACF profiles.'
SUP top ABSTRACT MAY ( racfOwner $ racfInstallationData $ racfDatasetModel ) )

objectclasses=( 1.3.18.0.2.6.58 NAME 'racfUser' DESC 'Represents a RACF USER Profile entry.' SUP racfbasecommon STRUCTURAL MUST ( racfid )

MAY ( racfAuthorizationDate $ racfAttributes $ racfPassword $ racfPasswordChangeDate $ racfPasswordInterval $ racfProgrammerName $ racfDefaultGroup $ racfLastAccess $ racfSecuritylabel $ racfSecurityCategoryList $ racfRevokeDate $ racfResumeDate $ racfLogonDays $ racfLogonTime $ racfClassName $ racfConnectGroupName $ racfConnectGroupNathority $ racfConnectGroupUACC $ racfSecurityLevel ) )

objectclasses=( 1.3.18.0.2.6.59 NAME 'racfGroup' DESC 'Represents a RACF GROUP Profile entry.' SUP racfbasecommon STRUCTURAL MUST ( racfid )
     MAY ( racfSuperiorGroup $ racfGroupNoTermUAC $ racfGroupUniversal $ racfSubGroupName $ racfGroupUserAccess $ racfGroupUserids ) )
objectclasses=( 1.3.18.0.2.6.60 NAME 'SAFDfpSegment' DESC 'Represents the SAF DFP portions of a RACF USER or GROUP profile.'
SUP top AUXILIARY MAY ( SAFDfpDataApplication $ SAFDfpDataClass $ SAFDfpManagementClass $ SAFDfpStorageClass ) )
objectclasses=( 1.3.18.0.2.6.61 NAME 'racfGroupOmvsSegment' DESC 'Represents the OS/390 OMVS Group information in a RACF GROUP profile.'
SUP top AUXILIARY MAY ( racfOmvsGroupId $ racfOmvsGroupIdKeyword ) ) objectclasses=( 1.3.18.0.2.6.62 NAME 'racfGroupOvmSegment' DESC 'Represents the OS/390 OVM Group information in a RACF GROUP profile.'
     SUP top AUXILIARY MAY ( racfOvmGroupId ) )
objectclasses= (1.3.18.0.2.6.63 NAME 'racfUserOmvsSegment' DESC 'Represents the OS/390 OMVS User information in a RACF USER profile.'
     SUP top AUXILIARY MAY ( racfOmvsUid $ racfOmvsHome $ racfOmvsInitialProgram $ racfOmvsMaximumAddressSpaceSize $ racfOmvsMaximumCPUTime $
     racfOmvsMaximumFilesPerProcess $ racfOmvsMaximumMemoryMapArea $ racfOmvsMaximumProcessesPerUID $ racfOmvsMaximumThreadsPerProcess $
     racfOmv sUidKeyword ) )
objectclasses=( 1.3.18.0.2.6.64 NAME 'racfUserOvmSegment' DESC 'Represents the OS/390 OVM User information protion of a RACF USER profile.'
SUP top AUXILIARY MAY ( racfOvmUid $ racfOvmHome $ racfOvmInitialProgram $ racfOvmFileSystemRoot ) )

objectclasses=( 1.3.18.0.2.6.65 NAME 'SAFTsoSegment' DESC 'Represents the OS/390 TSO information in a RACF USER profile.' SUP top

AUXILIARY MAY ( SAFAccountNumber $ SAFDestination $ SAFHoldClass $ SAFJobClass $ SAFMessageClass $ SAFDefaultLoginProc $ SAFLogonSize $
SAFMaximumRegionSize $ SAFDefaultSysoutClass $ SAFDefaultUnit $ SAFTsoSecurityLabel $ SAFDefaultCommand ) ) objectclasses=( 1.3.18.0.2.6.66 NAME 'racfCicsSegment' DESC 'Represents the OS/390 CICS information in a RACF USER profile.' SUP top
AUXILIARY MAY ( racfOperatorClass $ racfOperatorIdentification $ racfOperatorPriority $ racfOperatorReSignon $ racfTerminalTimeout ) ) objectclasses=( 1.3.18.0.2.6.67 NAME 'racfOperparmSegment' DESC 'Represents the OS/390 Operator parameters in a RACF USER profile.'
     SUP top AUXILIARY MAY ( racfStorageKeyword $ racfAuthKeyword $ racfMformKeyword $ racfLevelKeyword $ racfMonitorKeyword
     $ racfRoutcodeKeyword $
     racfLogCommandResponseKeyword \$ \ racfDDKeyword \$ \ racfEDMKeyword \$ \ racfKEYKeyword \$ \ racfCMDSYSKeyword \$ \ racfUDKeyword \$ \ racfDDKeyword \$ \ racfDD
     racfMscopeSystems $ racfAltGroupKeyword $ racfAutoKeyword ) )
objectclasses=( 1.3.18.0.2.6.68 NAME 'racfLanguageSegment' DESC 'Represents the OS/390 language information in a RACF USER profile.' SUP top
AUXILIARY MAY ( racfPrimaryLanguage $ racfSecondaryLanguage ) )
objectclasses=( 1.3.18.0.2.6.69 NAME 'racfWorkAttrSegment' DESC 'Represents the OS/390 work attributes information in a RACF USER profile.'
     SUP top AUXILIARY MAY ( racfWorkAttrUserName $ racfBuilding $ racfDepartment $ racfRoom $ racfAddressLine1 $ racfAddressLine2 $
racfAddress Line3 $ racfAddressLine4 $ racfWorkAttrAccountNumber ) )

objectclasses=( 1.3.18.0.2.6.70 NAME 'racfNetviewSegment' DESC 'Represents the OS/390 Netview information in a RACF USER profile.' SUP top

AUXILIARY MAY ( racfNetviewInitialCommand $ racfDefaultConsoleName $ racfCTLKeyword $ racfMSGRCVRKeyword $ racfNetviewOperatorClass $
racfDomains { racfNCMFADMKeyword ) } objectclasses=( 1.3.18.0.2.6.71 NAME 'racfDCESegment' DESC 'Represents the OS/390 DCE segment information in a RACF USER profile.'
     SUP top AUXILIARY MAY ( racfDCEAutoLogin $ racfDCEHomeCell $ racfDCEHomeCellUUID $ racfDCEPrincipal $ racfDCEUUID ) )
objectclasses=( 1.3.18.0.2.6.248 NAME 'racfLNotesSegment' DESC 'Represents the OS/390 LNOTES segment information in a RACF USER profile'
     SUP top AUXILIARY MAY ( racfLNotesShortName ) )
objectclasses=( 1.3.18.0.2.6.249 NAME 'racfNDSSegment' DESC 'Represents the OS/390 NDS segment information in a RACF USER profile'
SUP top AUXILIARY MAY ( racfNDSUserName ) ) objectclasses=( 1.3.18.0.2.6.260 NAME 'racfKerberosInfo' DESC 'Kerberos information for RACF'
SUP top AUXILIARY MAY ( krbPrincipalName \ maxTicketAge \ racfCurKeyVersion \ racfEncryptType ) )
objectclasses=( 1.3.18.0.2.6.259 NAME 'racfConnect' DESC 'RACF Connect' SUP top STRUCTURAL MUST ( racfUserId \ racfGroupId )
     MAY ( racfConnectAttributes $ racfConnectAuthDate $ racfConnectCount $ racfConnectGroupAuthority $
     racfConnectGroupUACC $ racfConnectLastConnect $ racfConnectOwner $ racfConnectResumeDate $ racfConnectRevokeDate ) )
objectclasses=( 1.3.18.0.2.6.267 NAME 'racfProxySegment' DESC 'RACF Proxy segment' SUP top AUXILIARY MAY ( racfLDAPBindDN $ racfLDAPBindPw $ racfLDAPHost ) ) objectclasses=( 1.3.18.0.2.6.447 NAME 'racfEIMSegment' DESC 'RACF EIM segment'
ODJECTCIASSES= ( 1.3.18.0.2.0.44 NAME 'PACTEIMSegment' DESC 'RACE ELM Segment' SUP top AUXILIARY MAY ( racfLDAPProf ) ) ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.3 DESC 'attribute type description' ) ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.12 DESC 'dn' ) ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.15 DESC 'directory string' )
 ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.16 DESC
                                                                                                             'dit content rule description'
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.17 DESC 'dit structure rule description' )
| Idapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.26 DESC 'ia5 string') | Idapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.26 DESC 'ia5 string') | Idapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.30 DESC 'matching rule description') | Idapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.31 DESC 'matching rule use description') | Idapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.35 DESC 'name form description') |
| ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.37 DESC 'object class description' ) | ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.38 DESC 'oid' )
ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.54 DESC 'ldap syntax description' ) ldapsyntaxes=( 1.3.6.1.4.1.1466.115.121.1.58 DESC 'substring assertion' )
ddapsyntaxes=(1.3.18.0.2.8.1 DESC 'ibm attribute type description')
matchingrules=(2.5.13.1 NAME 'distinguishedNameMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.12)
matchingrules=(2.5.13.2 NAME 'caseIgnoreMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)
matchingrules=(2.5.13.4 NAME 'caseIgnoreSubstringsMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.58)
matchingrules=( 1.3.6.1.4.1.1466.109.114.2 NAME 'caseIgnoreIA5Match' SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 )
matchingrules=( 2.5.13.5 NAME 'caseExactMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)
matchingrules=( 1.3.6.1.4.1.1466.109.114.1 NAME 'caseexactia5match' SYNTAX 1.3.6.1.4.1.1466.115.121.1.26)
matchingrules=( 2.5.13.14 NAME 'integerMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.27 )
matchingrules=( 2.5.13.29 NAME 'integerFirstComponentMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.27 )
matchingrules=( 2.5.13.0 NAME 'objectIdentifierMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 )
matchingrules=( 2.5.13.30 NAME 'objectIdentifierFirstComponentMatch' SYNTAX 1.3.6.1.4.1.1466.115.121.1.38 )
```

Appendix D. SPUFI files

This appendix shows the following SPUFI files:

- "The TDBMDB SPUFI file"
- "The TDBMINDX SPUFI file" on page 409

The TDBMDB SPUFI file

```
--* Licensed Materials - Property of IBM
--* 5694-A01
--* (C) Copyright IBM Corp. 2000
__**********************************
-- Use the following statements to create your LDAP Server DB2 database
-- and tablespaces in SPUFI. The database and tablespace names you
-- create will be used to update the database section of the LDAP
-- Server configuration file. You also need to make DB2 decisions,
-- in terms of buffer pool size selection for tablespaces and column
-- size selection, all of which will be directly related to the data that
-- will be stored in the database. See the instructions below for
-- more information.
__
__ *************
-- Database Name Information
__ **************
-- Change -DDDDDDDD- to the name of the LDAP database name you want to create.
-- Be sure this name is updated to match what is defined for databasename in
-- the server configuration file.
__ **************
-- DataBase Owner Information
__ **************
-- Change the -UUUUUUUU- to the MVS database owner id. This ID will be the
-- highlevel qualifier for the tables
__ *************
-- Tablespace Information
    *******************
    NOTE: Refer to the DB2 manuals for a complete listing of valid buffer
          pool names.
    *******************
-- Change the -AAAAAAAA to the LDAP entry tablespace name you want to create.
-- Change the -BBBB- to the buffer pool name for the LDAP entry tablespace.
      The size of the buffer pool can be determined with the formula:
--
       result = 62 bytes + <dn column trunc size (from below)> +
--
                   <maximum full size of a DN (from below)> +
        <size of entry data (which includes creator's DN and modifiers DN)>
      There is also a concept of a "spill over" table, where if the entry
      data does not fit into the row size, it will be broken up in order
      to fit into a row. Entry data may be spread across multiple rows
      if needed. So in the above formula, the <size of entry data>
      does not need to be the maximum size of the data, maybe the median
      size of the data would be a better choice. See the long entry
      tablespace description below.
```

TDBMDB SPUFI

```
The default suggested size is 4K.
-- Change the -CCCCCCC- to the LDAP long entry tablespace name you want to
-- create.
-- Change the -DDDD- to the buffer pool name for the LDAP long entry
       tablespace. The long entry table space will hold "spill over" rows
--
       for entry data that does not fit into the entry table tablespace.
--
      To minimize the number of spill over rows, choose a large buffer
--
      pool size.
      The default suggested size is 4K.
-- Change the -EEEEEEEE- to the LDAP long attribute tablespace name you want to
      create.
-- Change the -FFFF- to the buffer pool name for the LDAP long attribute
       tablespace. The long attribute table space will hold "spill over" rows
--
       for attribute data that does not fit into the entry table tablespace.
       To minimize the number of spill over rows, choose a large buffer
      pool size.
      The default suggested size is 4K.
-- Change the -GGGGGGGG- to the LDAP miscellaneous tablespace name you want
--
      to create.
__
-- Change the -HHHHHHHH- to the LDAP search tablespace name you want to create.
-- Change the -IIII- to the buffer pool name for the LDAP search tablespace.
      The size of the buffer pool can be determined with the simple formula:
--
--
     result = 16 bytes + <search column trunc size (from below)> +
              <maximum size of attribute value you would like to search for>
--
--
      The result value is the maximum number of bytes a row in the search
       table containing an attribute value will occupy. Choose a buffer pool
       size which will accomodate this size.
      The default suggested size is 4K.
-- Change the -JJJJJJJJ- to the LDAP replica tablespace name you want to
--
      create.
--
-- Change the -KKKKKKKK- to the LDAP descendants tablespace name you want
      to create.
__ **********
-- Column Size Selection Information
__ ****************
-- All searchable attributes of a given entry will be stored in two forms.
-- The first will be a truncated version, which will be used as part of
-- a DB2 index. The second version will be the entire attribute value,
-- potentially truncated by the buffer pool size you choose. The reason
-- two versions are stored is so that LDAP/DB2 can use indexes to increase
-- search performance. The reason we do not index the entire searchable
-- attribute value is because the cost (in terms of DASD) associated with
-- having indexes on a large column where there is a large amount of data.
-- The choice of the search column trunc size should take into account system
-- limits you may have (as described in the above), and should account
-- for the typical size of the attribute values that are stored in
-- LDAP. For example, if most of your data is only 20 bytes long,
-- choosing 20 for this trunc size would be wise.
-- Change -TTTT- to the search column trunc size you determine best fits your
-- attribute data.
```

```
The default suggested size is 32.
-- Another search performance enhancement is related to the DN attribute.
-- The DN attribute value is stored separately from the entry data to allow
-- a fast path lookup. It is also stored in two versions as well. The
-- reasons are similar to those mentioned above for the attribute column.
-- Since the DN data is stored in it's own column, you need to define the
-- maximum DN attribute value size here. You also need to choose a {\sf dn}
-- column trunc size that best fits your data.
-- Change -MMMM- to the dn trunc size you determine best fits your dn data.
      The default suggested size is 32.
-- Change -NNNN- to the maximum size of a DN.
--
      The default suggested size is 512.
--
__ **************
-- Storage Group Information
__ *************
-- Change the -SSSSSSS- to the storage group you want to contain the
-- LDAP DB2 tablespaces. Use SYSDEFLT to choose the default storage group.
    NOTE: The values provided below for PRIQTY and SECQTY probably need
          to be modified depending on the projected size of the
          Directory information to be stored.
-- *********************************
-- Use the following statements if you need to delete your LDAP Server DB2
-- database and tablespaces in SPUFI. You need to remove the '--'
-- from each line before you can run these statements.
-- Change the -AAAAAAAA to the LDAP entry tablespace name you want to delete.
-- Change the -CCCCCCC- to the LDAP long entry tablespace name you want to
                       delete.
-- Change the -EEEEEEEE- to the LDAP long attr tablespace name you want to
                       delete.
-- Change the -GGGGGGGG- to the LDAP miscellaneous tablespace name you want
                        to delete.
-- Change the -HHHHHHHH- to the LDAP search tablespace name you want to delete.
-- Change the -JJJJJJJJ- to the LDAP replica tablespace name you want to
                       delete.
-- Change the -KKKKKKKK- to the LDAP descendants tablespace name you want
                       to delete.
-- Change the -DDDDDDDD- to the LDAP database name you want to delete.
__ *********************************
-- DROP TABLESPACE -DDDDDDDD-.-AAAAAAA-;
-- DROP TABLESPACE -DDDDDDDD-.-CCCCCCCC-;
-- DROP TABLESPACE -DDDDDDDD-.-EEEEEEEE;
--DROP TABLESPACE -DDDDDDDD-.-GGGGGGGG-;
--DROP TABLESPACE -DDDDDDDD-.-HHHHHHHH-;
--DROP TABLESPACE -DDDDDDDD-.-JJJJJJJ-;
--DROP TABLESPACE -DDDDDDDD-.-KKKKKKKK-;
--DROP DATABASE -DDDDDDDD-;
--COMMIT;
__ *************
-- Create the LDAP database
__ ***************
CREATE DATABASE -DDDDDDDD- STOGROUP -SSSSSSSS-;
- *********
-- Create the LDAP entry tablespace
__ ******************
```

TDBMDB SPUFI

```
CREATE TABLESPACE -AAAAAAAA IN -DDDDDDDD-
      USING STOGROUP -SSSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL -BBBB-;
__ *************
-- Create the LDAP long entry tablespace
__ **********
CREATE TABLESPACE -CCCCCCC- IN -DDDDDDDD-
     USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL -DDDD-;
__ ************
-- Create the LDAP long attr tablespace
__ ***************
CREATE TABLESPACE -EEEEEEEE- IN -DDDDDDDD-
      USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL -FFFF-;
__ ***************
-- Create the LDAP 4K tablespace
__ ***************
CREATE TABLESPACE -GGGGGGGG- IN -DDDDDDDD-
     SEGSIZE 4
      USING STOGROUP -SSSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL BPO;
__ ****************
-- Create the LDAP search tablespace
__ *****************
CREATE TABLESPACE -HHHHHHHH- IN -DDDDDDDD-
     USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL -IIII-;
__ *****************
-- Create the LDAP replica tablespace
__ *****************
CREATE TABLESPACE -JJJJJJJJ- IN -DDDDDDDD-
      USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL -BBBB-;
__ ***************
-- Create the LDAP descendants tablespace
__ *****************
CREATE TABLESPACE -KKKKKKKK- IN -DDDDDDDD-
     USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 7200
     BUFFERPOOL BPO;
__ ***********
-- Create the DB2 tables
__ **********
__ *************
-- Create the DIR_ENTRY table
__ **************
CREATE TABLE -UUUUUUUU-.DIR ENTRY (
          DECIMAL(15,0)
DECIMAL(15,0),
                                     NOT NULL,
   PEID
                INTEGER,
INTEGER,
   ENTRY_SIZE
   LEVEL
                  DECIMAL(15,0),
   ACLSRC
   ACLPROP
                  CHAR(1),
   OWNSRC
                  DECIMAL(15,0),
   OWNPROP
                    CHAR(1),
   CREATE_TIMESTAMP
                    TIMESTAMP,
   MODIFY TIMESTAMP
                    TIMESTAMP,
   DN TRUNC
                    CHAR(-MMMM-) FOR BIT DATA,
                    VARCHAR(-NNNN-) FOR BIT DATA,
   DN
```

```
ENTRYDATA
                    LONG VARCHAR
                                    FOR BIT DATA,
   PRIMARY KEY( EID ) )
IN -DDDDDDD-.-AAAAAAA-;
__ **********
-- Create the DIR LONGENTRY table
__ ****************
CREATE TABLE -UUUUUUU-.DIR LONGENTRY (
   EID DECIMAL(15,0)
                                   NOT NULL,
   SEQ
                    INTEGER
                                   NOT NULL,
   ENTRYDATA
                    LONG VARCHAR
                                   FOR BIT DATA,
   PRIMARY KEY( EID, SEQ ) )
IN -DDDDDDDD-.-CCCCCCC-;
__ **************
-- Create the DIR LONGATTR table
__ ***************
CREATE TABLE -UUUUUUU-.DIR LONGATTR (
           DECIMAL(15,0)
                                   NOT NULL,
   EID
   ATTR ID
                    INTEGER
                                   NOT NULL,
                  INTEGER
   VALUENUM
                                   NOT NULL,
                   INTEGER
                                   NOT NULL,
   SEQ
                   LONG VARCHAR FOR BIT DATA,
   ATTRDATA
   PRIMARY KEY( EID, ATTR ID, VALUENUM, SEQ ) )
IN -DDDDDDDD-.-EEEEEEEE;
__ **************
-- Create the DIR MISC table
__ ***************
CREATE TABLE -UUUUUUUU-.DIR_MISC (
   NEXT_EID DECIMAL(15,0),
   NEXT ATTR ID
                   INTEGER,
               CHAR(10),
   DB VERSION
   DB CREATE VERSION CHAR(10) )
IN -DDDDDDDD-.-GGGGGGG-;
__ *************
-- Create the DIR CACHE table
__ **************
CREATE TABLE -UUUUUUU-.DIR CACHE (
   CACHE NAME CHAR(25)
                                   NOT NULL,
   MODIFY TIMESTAMP TIMESTAMP
                                   NOT NULL,
   PRIMARY KEY ( CACHE NAME, MODIFY TIMESTAMP ) )
IN -DDDDDDD-.-GGGGGGG-;
__ **************
-- Create the DIR ATTRID table
__ ***************
CREATE TABLE -UUUUUUU-.DIR_ATTRID (
            INTEGER,
   ATTR ID
   ATTR NOID
                   VARCHAR (200)
                                   NOT NULL,
   PRIMARY KEY( ATTR NOID ) )
IN -DDDDDDDD-.-GGGGGGGG-;
__ *************
-- Create the DIR DESC table
__ *************
CREATE TABLE -UUUUUUU-.DIR DESC (
                                   NOT NULL,
                    DECIMAL(15,0)
                    DECIMAL(15,0)
                                   NOT NULL,
   PRIMARY KEY( DEID, AEID ) )
IN -DDDDDDDD-.-KKKKKKK-;
__ *************
-- Create the DIR SEARCH table
__ **************
CREATE TABLE -UUUUUUU-.DIR SEARCH (
```

TDBMDB SPUFI

```
DECIMAL(15,0)
                                     NOT NULL,
   EID
   ATTR ID
                     INTEGER
                                     NOT NULL,
   VALUE
                     CHAR(-TTTT-)
                                     FOR BIT DATA,
                                     FOR BIT DATA )
                     LONG VARCHAR
   LVALUE
IN -DDDDDDDD-.-HHHHHHHH-;
__ ****************
-- Create the DIR REGISTER table
__ **************
CREATE TABLE -UUUUUUU-.DIR REGISTER (
                     INTEGER
                                     NOT NULL,
   SRV
                     VARCHAR (125)
                                     NOT NULL,
   PRIMARY KEY( ID, SRV ) )
IN -DDDDDDDD-.-GGGGGGG-;
__ ***************
-- Create the DIR PROGRESS table
__ ***************
CREATE TABLE -UUUUUUU-.DIR PROGRESS (
   ID
                     INTEGER
                                     NOT NULL,
   PRG
                     VARCHAR (125)
                                     NOT NULL,
   SRV
                     VARCHAR (125)
                                     NOT NULL,
   PRIMARY KEY( ID, PRG, SRV ) )
IN -DDDDDDDD-.-GGGGGGG-;
__ **************
-- Create the DIR CHANGE table
__ **************
CREATE TABLE -UUUUUUU-.DIR CHANGE (
                                     NOT NULL,
   ΙD
                    INTEGER
   TYPE
                     INTEGER
                                     NOT NULL,
   LONGENTRY_SIZE
                     INTEGER,
                     VARCHAR (-NNNN-)
                                     NOT NULL,
   DIN
                     LONG VARCHAR
                                     NOT NULL,
   PRIMARY KEY( ID ) )
IN -DDDDDDDD-.-JJJJJJJ-;
__ ***************
-- Create the DIR LONGCHANGE table
__ ***************
CREATE TABLE -UUUUUUU-.DIR LONGCHANGE (
                    INTEGER
                                     NOT NULL,
   SEQ
                     INTEGER
                                     NOT NULL,
   LDIF
                     LONG VARCHAR,
   PRIMARY KEY( ID, SEQ ) )
IN -DDDDDDDD-.-JJJJJJJ-;
__ ****************
-- Commit all the above SQL statements
__ *****************
COMMIT;
```

The TDBMINDX SPUFI file

```
__***********************
--*
--* Licensed Materials - Property of IBM
--* 5694-A01
--* (C) Copyright IBM Corp. 2000
__*********************************
-- Use the following statements to create your LDAP Server DB2
-- indexes in SPUFI. See the instructions below for more information.
__ **************
-- DataBase Owner Information
__ *************
-- Change the -UUUUUUUU- to the MVS database owner id. This ID will be the
-- highlevel qualifier for the tables. This value should correspond
-- with the value chosen in the LDAP Server DB2 database and tablespace
-- SPUFI script.
__ **************
-- Storage Group Information
__ ************
-- Change the -SSSSSSSS- to the storage group you want to contain the
-- LDAP DB2 indexes. Use SYSDEFLT to choose the default storage group.
    NOTE: The values provided below for PRIQTY and SECQTY probably need
         to be modified depending on the projected size of the
         Directory information to be stored.
__ **************
-- Miscellaneous Information
__ **************
-- All indexes have been defined DEFER YES, which means they need to be
-- recovered at some point. It is suggested to do the recovery after
-- the database has been populated for databases with large amounts of
-- data. Use of this option is strictly optional though.
-- To not use the DEFER YES option, simply remove DEFER YES globablly.
__ **************
-- Create the DIR ENTRY indexes
__ **************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR ENTRYX0 ON -UUUUUUUU-.DIR ENTRY( EID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
CREATE INDEX -UUUUUUU-.DIR ENTRYX1 ON -UUUUUUUU-.DIR ENTRY( PEID, EID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
CREATE INDEX -UUUUUUU-.DIR ENTRYX2 ON -UUUUUUUU-.DIR ENTRY( EID, DN TRUNC )
   USING STOGROUP -SSSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
CREATE INDEX -UUUUUUU-.DIR ENTRYX3 ON -UUUUUUUU-.DIR ENTRY( DN TRUNC, EID )
   USING STOGROUP -SSSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ *****************
-- Create the DIR LONGENTRY indexes
__ **********
CREATE UNIQUE INDEX -UUUUUUUU-.DIR LONGENTRYX1
   ON -UUUUUUU-.DIR LONGENTRY( EID, SEQ )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
-- **********
-- Create the DIR LONGATTR indexes
```

TDBMINDX SPUFI

```
__ **************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR LONGATTRX1
   ON -UUUUUUU-.DIR LONGATTR( EID, ATTR ID, VALUENUM, SEQ )
   USING STOGROUP -SSSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ ***************
-- Create the DIR CACHE indexes
__ **************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR CACHEX1
   ON -UUUUUUU-.DIR CACHE( CACHE NAME, MODIFY TIMESTAMP )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ **************
-- Create the DIR ATTRID indexes
__ ****************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR ATTRIDX1
   ON -UUUUUUU-.DIR_ATTRID( ATTR_NOID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ **************
-- Create the DIR DESC indexes
__ ***************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR DESCX1
   ON -UUUUUUU-.DIR_DESC( DEID, AEID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ ***************
-- Create the DIR_SEARCH indexes
__ ****************
CREATE INDEX -UUUUUUU-.DIR SEARCHX1
   ON -UUUUUUU-.DIR_SEARCH( ATTR_ID, VALUE, EID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
CREATE INDEX -UUUUUUUU-.DIR SEARCHX2
   ON -UUUUUUU-.DIR SEARCH( EID, ATTR ID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160 CLUSTER
   DEFER YES;
__ ****************
-- Create the DIR REGISTER indexes
__ ***************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR REGISTERX1
   ON -UUUUUUU-.DIR REGISTER( ID, SRV )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ ***************
-- Create the DIR PROGRESS indexes
__ ****************
CREATE UNIQUE INDEX -UUUUUUUU-.DIR PROGRESSX1
   ON -UUUUUUU-.DIR PROGRESS( ID, PRG, SRV )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ ***************
-- Create the DIR CHANGE indexes
__ ****************
CREATE UNIQUE INDEX -UUUUUUU-.DIR_CHANGEX1 ON -UUUUUUUU-.DIR_CHANGE( ID )
   USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160
   DEFER YES;
__ ******************
```

TDBMINDX SPUFI

-- Create the DIR_LONGCHANGE indexes __ *************** CREATE UNIQUE INDEX -UUUUUUUU-.DIR_LONGCHANGEX1 ON -UUUUUUU-.DIR_LONGCHANGE(TD, SEQ) USING STOGROUP -SSSSSSS- PRIQTY 14400 SECQTY 2160 DEFER YES; __ ****************** -- Commit all the above SQL statements __ ***************** COMMIT;

TDBMINDX SPUFI

Appendix E. Sample JCL

This appendix shows the following JCL samples:

- · "Sample JCL for the LDAP server"
- "Sample JCL for Idif2tdbm" on page 415
- "Sample JCL for tdbm2ldif" on page 417

Sample JCL for the LDAP server

Following is the sample JCL provided for the LDAP server (LDAPSRV PROC JCL).

```
//*
//* Licensed Materials - Property of IBM
//* 5694-A01
//* (C) Copyright IBM Corp. 1997, 1999
//*
//* Procedure for starting the LDAPSRV server
//* To start server using configuration file
//* /etc/ldap/slapd.conf specify:
//* s ldapsrv
//*
//* To start server using alternate configuration file or
//* other parameters specify:
//* s ldapsrv,parms='options'
//* where options can be:
//*
         -f filename # alternate configuration file
//*
         -d level # debug level (65535 turns on all debugs)
         -p portno  # non-secure port number (deprecated)
-s portno  # secure port number (deprecated)
//*
//*
          -1 listenURL # ldap URL
//* An alternative to the -f option is to define a CONFIG DD.
//* The remaining options are optional. If not set, message/debug
//* levels are set to 0, non-secure port number will be 389, and
//* secure port number will be 636. NOTE: use of these low port
//* numbers will require that the LDAPSRV server run under a userid
//* that has OpenEdition UID 0.
//*
//* In the JCL below, GLDHLQ refers to the high level qualifier
//* that was used to install the LDAP Server datasets. This will
//* have to be customized for the installation.
//*
//LDAPSRV PROC REGSIZE=64M,
//*----
//* CUSTOMIZABLE SYMBOLIC PARAMETERS
// PARMS='',
// GLDHLQ='XXXXXX',
// OUTCLASS='A'
//GO EXEC PGM=GLDSLAPD, REGION=&REGSIZE, TIME=1440,
      PARM=('/&PARMS >DD:SLAPDOUT 2>&1')
//*-----
//* STEPLIB must be customized based on install HLQ.
//STEPLIB DD DSN=&GLDHLQ..SGLDLNK,DISP=SHR
//*-----
```

Sample JCL (LDAP Server)

```
//* Fill in and uncomment the following DD if the libraries for DB2
//* are not in LINKLST or LPA on the system. Modify <DB2HLQ>
//* to be the high level qualifier of the DB2 installation
//* on the system.
//* DD DSN=<DB2HLQ>.SDSNLOAD,DISP=SHR
//*-----
//* CONFIG can be used to specify the LDAP server config file.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*CONFIG DD DSN=<CONFIG.FILE.DATASET>,DISP=SHR
//*-----
//* ENVVAR can be used to specify any environment variables
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*-----
//*ENVVAR DD DSN=<ENVVAR.FILE.DATASET>,DISP=SHR
//*-----
//* DSNAOINI can be used to specify the dsnaoini file required by DB2.
//* Alternatively, this data set can also be specified in the
//* configuration file.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*DSNAOINI DD DSN=<DSNAOINI.DATASET>,DISP=SHR
//SLAPDOUT DD SYSOUT=&OUTCLASS
//SYSOUT DD SYSOUT=&OUTCLASS
//SYSUDUMP DD SYSOUT=&OUTCLASS
//CEEDUMP DD SYSOUT=&OUTCLASS
```

Sample JCL for Idif2tdbm

```
//LDF2TDBM JOB MSGCLASS=H
//*
//* Licensed Materials - Property of IBM
//* 5694-A01
//* (C) Copyright IBM Corp. 2000
//*
//* Procedure for running LDF2TDBM from batch
//*
//* LDF2TDBM has several required parameters and many optional ones.
//* You must update the DSHLQ, STEPS, and INFILE substitution
//* parameters below to specify the required parameters. Optional
//* parameters can be put in the PARMS substitution parameter. The
//* required parameters are:
//*
//*
                     # high-level qualifiers of the output datasets
//* <-c | -p | -1>
                   # one or more processing steps to perform
//*
                      (check, prepare, and/or load)
//*
     <-i | -e> filename # name of input file and/or name of a file
//*
                       containing a list of input files
//*
//* One optional parameter has already been added in PARMS
//* -a <yes | no> # whether to continue despite warning messages
//* Do not remove the -a option because processing will be suspended
//* if a warning message is issued.
//*
//* Other optional parameters include:
//*
//*
                     # alternate configuration file
     -f filename
//*
                     # name of schema input file
    -s filename
//*
//* See the LDAP Server Administration and Usage Guide for a full
//* description of all the required and optional parameters.
//* An alternative to the -f option is to define a CONFIG DD
//* in the PROC. If the -f option is not specified and the CONFIG DD //* card is not specified, the HFS file '/etc/slapd/slapd.conf' will be
//* used for the config file.
//* In the JCL below, GLDHLQ refers to the high level qualifier
//* that was used to install the LDAP Server datasets. This will
//* have to be customized for the installation.
//LDF2TDBM PROC REGSIZE=0M.
//* CUSTOMIZABLE SYMBOLIC PARAMETERS
//* Customize the required parameters here and add optional parameters
//* for the desired behavior. DSHLQ, STEPS, and INFILE refer to the
//* three required parameters mentioned above. In INFILE, you can
//* specify multiple input files and/or use the -e option to specify
//* a file containing a list of input files.
//*-----
// DSHLQ='-o YYYYY',
// STEPS='-cpl',
// INFILE='-i ZZZZZ',
// PARMS='-a WWW',
// GLDHLO='XXXXXXX'.
// OUTCLASS='A'
//LDF2TDBM EXEC PGM=GLDBKLD, REGION=&REGSIZE,
         PARM=('/&DSHLQ &STEPS &INFILE &PARMS')
```

Sample JCL (Idif2tdbm)

```
//* STEPLIB must be customized based on install HLQ.
//*-----
//STEPLIB DD DSN=&GLDHLQ..SGLDLNK,DISP=SHR
//*-----
//* Fill in and uncomment the following DD if the libraries for DB2
//* are not in LINKLST or LPA on the system. Modify <DB2HLQ>
//* to be the high level qualifier of the DB2 installation
//* on the system.
//*-----
//* DD DSN=<DB2HLQ>.SDSNLOAD,DISP=SHR
//*-----
//* DSNAOINI can be used to specify the DSNAOINI file needed for DB2.
//* If this DD is not used, the DSNAOINI setting must be in the
//* LDAP server configuration file.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*-----
//*DSNAOINI DD DSN=<DB2.DSNAOINI.DATASET>,DISP=SHR
//*-----
//* CONFIG can be used to specify the LDAP server config file.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*-----
//*CONFIG DD DSN=<CONFIG.FILE.DATASET>,DISP=SHR
//*-----
//* ENVVAR can be used to specify any environment variables
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*ENVVAR DD DSN=<ENVVAR.FILE.DATASET>,DISP=SHR
//*-----
//SYSPRINT DD SYSOUT=&OUTCLASS
//CEEDUMP DD SYSOUT=&OUTCLASS
//SYSERR DD SYSOUT=&OUTCLASS
//STDOUT DD SYSOUT=&OUTCLASS
// PEND
//*-----
//GO EXEC LDF2TDBM
```

Sample JCL for tdbm2ldif

```
//TDBM2LDF JOB MSGCLASS=H
//*
//* Licensed Materials - Property of IBM
//* 5694-A01
//* (C) Copyright IBM Corp. 2000
//*
//*
//* Procedure for running TDBM2LDF from batch
//*
//* To start TDBM2LDF using configuration file
//* /etc/ldap/slapd.conf just submit the job.
//* To start TDBM2LDF using an alternate configuration file or
//* other parameters specify one or more of the following
//* options in the PARMS substitution parameter in the
//* EXEC line below:
//*
//* The options can be:
//*
          -f filename # alternate configuration file
//*
           -o filename # file in which to store the output
//*
           -s subtreeDN # root of the subtree to extract
//*
                     # output userpassword in tagged format
//*
//* An alternative to the -f option is to define a CONFIG DD
//* in the PROC.
//*
//* An alternative to the -o option is to define the SYSPRINT DD
//* card to a dataset that will store the output.
//* The -f and -o, -s, and -t options are optional. If the -f option
//* is not specified and the CONFIG DD card is not specified, the
//* HFS file '/etc/slapd/slapd.conf' will be used for the config file.
//* If the -o option is not specified, then the SYSPRINT DD card is
//* is used for the output from the program. If the -s option is
//* not specified, all data held by the LDAP server will be unloaded.
//*
//* In the JCL below, GLDHLQ refers to the high level qualifier
//* that was used to install the LDAP Server datasets. This will
//* have to be customized for the installation.
//*
//TDBM2LDF PROC REGSIZE=0M,
//* CUSTOMIZABLE SYMBOLIC PARAMETERS
//* Customize the parameters here for desired behavior.
// PARMS='',
// GLDHLQ='XXXXXX',
// OUTCLASS='A'
//TDBM2LDF EXEC PGM=GLDUNLD, REGION=&REGSIZE,
   PARM=('/&PARMS')
//*----
//* STEPLIB must be customized based on install HLQ.
//STEPLIB DD DSN=&GLDHLQ..SGLDLNK,DISP=SHR
//*----
//* Fill in and uncomment the following DD if the libraries for DB2
//* are not in LINKLST or LPA on the system. Modify <DB2HLQ>
//* to be the high level qualifier of the DB2 installation
//* on the system.
//*-----
```

Sample JCL (tdbm2ldif)

```
DD DSN=<DB2HLQ>.SDSNLOAD,DISP=SHR
//*----
//* DSNAOINI can be used to specify the DSNAOINI file needed for DB2.
//* If this DD is not used, the DSNAOINI setting must be in the
//* LDAP server configuration file.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*-----
//*DSNAOINI DD DSN=<DB2.DSNAOINI.DATASET>,DISP=SHR
//* CONFIG can be used to specify the LDAP server config file.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*CONFIG DD DSN=<CONFIG.FILE.DATASET>,DISP=SHR
//*-----
//* ENVVAR can be used to specify any environment variables
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*ENVVAR DD DSN=<ENVVAR.FILE.DATASET>,DISP=SHR
//*-----
//* SYSPRINT can be used to specify the location for output
//* from this command.
//* If this DD is used, the name of the dataset must be customized
//* for the installation.
//*-----
//*SYSPRINT DD DSN=<OUTPUT.LDIF.DATASET>,DISP=NEW
//SYSPRINT DD SYSOUT=&OUTCLASS
//CEEDUMP DD SYSOUT=&OUTCLASS
//SYSERR DD SYSOUT=&OUTCLASS
//STDOUT DD SYSOUT=&OUTCLASS
// PEND
//*-----
//GO EXEC TDBM2LDF
```

Sample JCL for Idapadduuids

```
//GLDADDUU JOB MSGCLASS=H
//*
//* Licensed Materials - Property of IBM
//* 5694-A01
//* (C) Copyright IBM Corp. 2002
//*
//*
//* Procedure for running ldapadduuids from batch
//*
//* To start specify one or more of the following
//* options in the PARMS substitution parameter in the
//* EXEC line below:
//*
//* The options can be:
//*
     -S mechanism select SASL bind mechanism (supported mechanisms
//*
     -b basedn base dn for search.
//*
     -s scope one of base, one, or sub (search scope)
//*
     -1 time lim time limit (in seconds) for search
//*
     -z size lim size limit (in entries) for search
     -D binddn
//*
                bind dn
//*
     -w passwd
                 bind passwd
     -w pus
-h host
//*
                ldap server
//*
     -p port
                port on ldap server
//*
                use a secure ldap connection for search
     -Z
//*
     -K keyfile file to use for keys/certificates
//*
                keyfile password
     -P key pw
//*
                 Certificate label in keyfile
     -N key dn
//*
     -m realm
                 Mandatory Authentication realm
//*
     -U username Mandatory Authentication username (uid)
//*
     -n don't do operation
//*
     -d debug level
//*
//*
     followed by an optional search filter.
//*
//* All options are optional, but program will not run correctly
//*
     without proper settings for -D, -w and -b.
//*
//* If the -h option is not specified, localhost is assumed.
//* If the -p option is not specified, port 389 is assumed.
//* If the -b option is not specified, the environment variable
//*
     LDAP_BASEDN will be examined. If no base is found,
//*
     the program will fail.
//*
//* Specify the -Z and -K options to use an SSL connection to
//* the server for this run of the ldapadduuids program. It may also
//* be necessary to specify -P and -N depending on the setup of
//* the keyfile specified in -K.
//*
//* In the JCL below, GLDHLQ refers to the high level qualifier
//* that was used to install the LDAP Server datasets. This will
//* have to be customized for the installation.
//*
//GLDADDU PROC REGSIZE=2048K,
//*-----
//* CUSTOMIZABLE SYMBOLIC PARAMETERS
//* Customize the parameters here for desired behavior.
//*-----
// PARMS='-D "CN=xxxxxx" -w xxxxxx -b "o=xxxxx,c=xxxxx"',
// GLDHLQ='xxxxxx',
// OUTCLASS='A'
```

Sample JCL (tdbm2ldif)

```
//GLDADDU EXEC PGM=GLDADDUU, REGION=&REGSIZE.
// PARM=('ENVAR("_CEE_ENVFILE=DD:ENVVARS")/&PARMS')
//*------
//* STEPLIB must be customized based on install HLQ.
//*-----
//STEPLIB DD DSN=&GLDHLQ..SGLDLNK,DISP=SHR
//*-----
//* If using SSL, the following DD must be uncommented. The
//* dataset name where SystemSSL was installed must be
//* modified for your installation.
//* DD DSN=<SystemSSL HLQ>.SGSKLOAD,DISP=SHR
//*----
//SYSPRINT DD SYSOUT=&OUTCLASS
//CEEDUMP DD SYSOUT=&OUTCLASS
//SYSERR DD SYSOUT=&OUTCLASS
//STDOUT DD SYSOUT=&OUTCLASS
//*-----
//* To customize the environment variables, copy the HFS file noted
//* below to another HFS file, customize the environment variable
//* settings and change the PATH statement below to indicate the new
//* file name.
//*-----
//ENVVARS DD PATH='/usr/lpp/ldap/etc/slapd.envvars',PATHOPTS=ORDONLY
// PEND
//*-----
//GO EXEC GLDADDU
```

Appendix F. Sample LDIF input file

The following sample LDIF input file can be found in the /usr/lpp/ldap/examples/sample_server directory and is called sample.ldif.

```
dn: o=Your Company
objectclass: top
objectclass: organization
o: Your Company
dn: cn=LDAP Administrator, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: LDAP Administrator
sn: Administrator
userPassword: secret
dn: ou=Home Town, o=Your Company
objectclass: top
objectclass: organizationalUnit
ou: Home Town
seealso: cn=Linda Carlesberg, ou=Home Town, o=Your Company
dn: ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: organizationalUnit
ou: In Flight Systems
description: main product: Course Maker
businessCategory: aircraft
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
dn: ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: organizationalUnit
ou: Home Entertainment
description: main product: TV Connection
businessCategory: Home Entertainment
dn: ou=Groups, o=Your Company
objectclass: top
objectclass: organizationalUnit
ou: Groups
dn: cn=Bowling team, ou=Groups, o=Your Company
objectclass: top
objectclass: groupOfNames
description: IBM Home Town Bowling Team
cn: Bowling team
owner: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
member: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
member: cn=Michael Campbell+postalcode=4609, ou=Widget Division, ou=Home Town, o=Your Company
member: cn=Eddie Catu, ou=In Flight Systems, ou=Home Town, o=Your Company
member: cn=Melinda Charles, ou=In Flight Systems, ou=Home Town, o=Your Company
member: cn=Al Edwards, ou=Widget Division, ou=Home Town, o=Your Company
dn: ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: organizationalUnit
ou: Widget Division
description: main product: Orange Widget Delux
businessCategory: home entertainment
dn: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
```

```
objectclass: person
objectclass: organizationalPerson
cn: Mary Burnnet
sn: Burnnet
telephonenumber: 1-812-855-5923
internationaliSDNNumber: 755-5923
facsimiletelephonenumber: 1-812-855-5923
title: ISO Deputy, Qual. Tech
postalcode: 1515
seealso: cn=Linda Carlesberg, ou=Home Town, o=Your Company
dn: cn=David Campbell, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: David Campbell
sn: Campbell
telephonenumber: 1-812-855-7509
internationaliSDNNumber: 755-7509
title: Mfg. Assembly
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
postalcode: 1514
dn: cn=James Campbell, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: James Campbell
sn: Campbell
telephonenumber: 1-812-855-8861
internationaliSDNNumber: 755-8861
facsimiletelephonenumber: 1-812-855-5237
title: Home Town Adaptive Technology Center Accessibility
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
postalcode: 4503
telexnumber: 1-812-343-7700
facsimiletelephonenumber: 755-5237
dn: cn=Michael Campbell, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Michael Campbell
sn: Campbell
telephonenumber: 1-812-855-5838
internationaliSDNNumber: 755-5838
postalcode: 4681
dn: cn=Michael Campbell+postalcode=4609, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Michael Campbell
sn: Campbell
telephonenumber: 1-812-855-7743
title: Drill Department
postalcode: 4609
dn: cn=Bob Campbell, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Bob Campbell
sn: Campbell
telephonenumber: 1-812-855-8541
internationaliSDNNumber: 755-8541
title: Mechanical Ana. Thermal
```

```
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
postalcode: 4502
dn: cn=Bonnie Daniel, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Bonnie Daniel
sn: Daniel
telephonenumber: 1-812-855-7453
internationaliSDNNumber: 755-7453
title: RISC Manufacturing
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
postalcode: 1515
dn: cn=Brenda England, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Brenda England
sn: England
telephonenumber: 1-812-855-5231
internationaliSDNNumber: 755-5231
facsimiletelephonenumber: 1-812-855-5237
title: Assistant to Dr. Campbell
postalcode: 4503
facsimiletelephonenumber: 755-5237
dn: cn=David Delbert, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: David Delbert
sn: Delbert
telephonenumber: 1-812-855-8504
internationaliSDNNumber: 523-8504
facsimiletelephonenumber: 1-812-855-6040
title: SWAT OS/2 Analyst
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
postalcode: 2901
telexnumber: 1-800-546-4646
facsimiletelephonenumber: 755-6040
dn: cn=Al Edwards, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Al Edwards
sn: Edwards
telephonenumber: 1-812-855-1370
internationaliSDNNumber: 755-1370
facsimiletelephonenumber: 1-812-855-5004
title: Site Occupancy Planner
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
postalcode: 4713
dn: cn=Arthur Edwards, ou=Widget Division, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Arthur Edwards
sn: Edwards
telephonenumber: 1-812-855-8474
internationaliSDNNumber: 523-8474
facsimiletelephonenumber: 1-812-855-6040
title: PSP Enterprise Customer Support OS/2 SWAT Team
seealso: cn=Mary Burnnet, ou=Widget Division, ou=Home Town, o=Your Company
```

postalcode: 2901 telexnumber: 1-800-546-4646 facsimiletelephonenumber: 755-6040 dn: cn=Curtis Edwards Jr, ou=Widget Division, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Curtis Edwards Jr sn: Edwards telephonenumber: 1-812-855-8053 internationaliSDNNumber: 755-8053 facsimiletelephonenumber: 1-812-855-7101 title: EMC TEST postalcode: 4502 facsimiletelephonenumber: 755-7101 dn: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Cynthia Flowers sn: Flowers telephonenumber: 1-812-855-8609 internationaliSDNNumber: 755-8609 facsimiletelephonenumber: 1-812-855-8712 title: Software Contracts postalcode: 1725 facsimiletelephonenumber: 755-8712 seealso: cn=Linda Carlesberg, ou=Home Town, o=Your Company dn: cn=Doug Edwards, ou=Home Entertainment, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Doug Edwards sn: Edwards telephonenumber: 1-812-855-8386 internationaliSDNNumber: 755-8386 facsimiletelephonenumber: 1-812-855-8199 title: Panel Finance / Accounting seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company postalcode: 4604 facsimiletelephonenumber: 755-8199 dn: cn=Jeffrey James, ou=Home Entertainment, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Jeffrey James sn: James telephonenumber: 1-812-855-7551 internationaliSDNNumber: 755-7551 facsimiletelephonenumber: 1-812-855-7193 title: programmer seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company postalcode: 1033 facsimiletelephonenumber: 755-7193 dn: cn=Ron Edwards, ou=Home Entertainment, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Ron Edwards sn: Edwards telephonenumber: 1-812-855-4021 internationaliSDNNumber: 755-4021

```
facsimiletelephonenumber: 1-812-855-5454
title: DEPT TECH
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 4601
telexnumber: 1-812-474-3783
dn: cn=Jerry Chevy, ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Jerry Chevy
sn: Chevy
telephonenumber: 1-812-855-7562
internationaliSDNNumber: 755-7562
facsimiletelephonenumber: 1-812-855-5004
title: SITE STRATEGIC PLANNER
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 4713
facsimiletelephonenumber: 755-5004
dn: cn=Marvin McGee, ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Marvin McGee
sn: McGee
telephonenumber: 1-812-855-9797
internationaliSDNNumber: 755-9797
facsimiletelephonenumber: 1-812-855-5004
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 4713
facsimiletelephonenumber: 755-5004
dn: cn=Marshall Riely, ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Marshall Riely
sn: Riely
telephonenumber: 1-812-855-7218
internationaliSDNNumber: 755-7218
title: auto. equip. maint. spec.
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 4601
telexnumber: 1-812-480-7509
dn: cn=James Giliam, ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: James Giliam
sn: Giliam
telephonenumber: 1-812-855-5386
internationaliSDNNumber: 755-5386
facsimiletelephonenumber: 1-812-855-5824
title: Project Management
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 9635
dn: cn=Al Garcia, ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Al Garcia
sn: Garcia
telephonenumber: 1-812-855-7579
internationaliSDNNumber: 755-7095
```

```
title: LEAD TA / MAINTENANCE
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 1377
dn: cn=Ben Garcia Jr, ou=Home Entertainment, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Ben Garcia Jr
sn: Garcia
telephonenumber: 1-812-855-3674
internationaliSDNNumber: 523-3674
facsimiletelephonenumber: 1-812-855-1077
title: OS/2 LAN Server Support
seealso: cn=Cynthia Flowers, ou=Home Entertainment, ou=Home Town, o=Your Company
postalcode: 2901
telexnumber: 1-812-474-3111
facsimiletelephonenumber: 755-1077
dn: cn=Becky Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Becky Garcia
sn: Garcia
telephonenumber: 1-812-855-5366
internationaliSDNNumber: 755-5366
facsimiletelephonenumber: 1-812-855-7961
title: Flight Manager in Professional Certification
postalcode: 9635
facsimiletelephonenumber: 755-7961
seealso: cn=Linda Carlesberg, ou=Home Town, o=Your Company
dn: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Maria Garcia
sn: Garcia
telephonenumber: 1-812-855-8717
internationaliSDNNumber: 755-8717
title: SUPPLEMENTAL
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 4633
dn: cn=Bob Garcia, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Bob Garcia
sn: Garcia
telephonenumber: 1-812-855-4553
internationaliSDNNumber: 755-4553
title: Preload Test
seealso: cn=Linda Carlesberg, ou=Home Town, o=Your Company
postalcode: 9340
dn: cn=Ricardo Garcia, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Ricardo Garcia
sn: Garcia
telephonenumber: 1-812-855-8278
internationaliSDNNumber: 755-8278
postalcode: 1365
```

```
dn: cn=Amy Nguyen, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Amy Nguyen
sn: Nguyen
telephonenumber: 1-812-855-7189
internationaliSDNNumber: 755-7189
facsimiletelephonenumber: 1-812-855-8199
title: Coop Dept 32E
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 4604
dn: cn=James Nguyen, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: James Nguyen
sn: Nguyen
telephonenumber: 1-812-855-4156
internationaliSDNNumber: 755-4156
title: AIX Network Device Drivers
postalcode: 9551
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
dn: cn=Henry Nguyen, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Henry Nguyen
sn: Nguyen
telephonenumber: 1-812-855-4028
internationaliSDNNumber: 755-4028
facsimiletelephonenumber: 1-812-855-9087
title: AIX Support
postalcode: 9551
facsimiletelephonenumber: 755-9087
dn: cn=Kyle Nguyen, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Kyle Nguyen
sn: Nguyen
telephonenumber: 1-812-855-8974
internationaliSDNNumber: 755-8974
title: System Support
seealso: cn=Linda Carlesberg, ou=Home Town, o=Your Company
postalcode: 9810
telexnumber: 1-812-397-1205
dn: cn=Wayne Nguyen, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Wayne Nguyen
sn: Nguyen
telephonenumber: 1-812-855-5052
internationaliSDNNumber: 755-5052
title: Object technology consultant
postalcode: 1003
dn: cn=Jason Li, ou=In Flight Systems, ou=Home Town, o=Your Company
bjectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Jason Li
```

```
sn: Li
telephonenumber: 1-812-855-1466
internationaliSDNNumber: 755-1466
facsimiletelephonenumber: 1-812-855-3882
title: Internet Solutions; HA
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 9584
facsimiletelephonenumber: 755-3882
dn: cn=Melinda Charles, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Melinda Charles
sn: Charles
telephonenumber: 1-812-855-5489
internationaliSDNNumber: 755-5489
facsimiletelephonenumber: 1-812-855-7670
title: Integrated Procurement Solutions Home Town
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 1109
facsimiletelephonenumber: 755-7670
dn: cn=Bill Keller Jr., ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Bill Keller Jr.
sn: Keller
telephonenumber: 1-812-855-5245
internationaliSDNNumber: 755-5245
facsimiletelephonenumber: 1-812-855-8138
title: Returned parts inventory control
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 1505
telexnumber: 1-800-563-7138
facsimiletelephonenumber: 755-8138
dn: cn=Cynthia Smith, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Cynthia Smith
sn: Smith
telephonenumber: 1-812-855-8301
internationaliSDNNumber: 755-8301
facsimiletelephonenumber: 1-812-855-6074
title: Electrical Analysis
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 9812
facsimiletelephonenumber: 755-6074
dn: cn=Donald Sinclar, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Donald Sinclar
sn: Sinclar
telephonenumber: 1-812-855-8840
internationaliSDNNumber: 755-8840
facsimiletelephonenumber: 1-812-855-8138
title: Mgr. Returned RISC Mach. / Recon / Scrap
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
postalcode: 1514
facsimiletelephonenumber: 755-8138
dn: cn=Ben Catu, ou=In Flight Systems, ou=Home Town, o=Your Company
```

```
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Ben Catu
sn: Catu
telephonenumber: 1-812-855-6218
internationaliSDNNumber: 755-6218
postalcode: 9811
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
dn: cn=Eddie Catu, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
description: This entry has every attribute for this class
sn: Catu
description: He puts it all together
destinationindicator: final assembly
facsimiletelephonenumber: 1-812-855-8985
internationaliSDNNumber: 755-8498
1: North America
physicaldeliveryofficename: doc 17
postofficebox: 1420
postaladdress: 2183 Tamil In. Home Town TX 78659
postalcode: 1800
preferreddeliveryMethod: UPS
registeredaddress: Catu.ring2.austin.ibm.com
st: TX
street: 2183 Tamil ln.
telephonenumber: 1-812-855-8498
teletexterminalidentifier: 755-8498
telexnumber: 755-8498
title: Assembly
userpassword:: Y2hhbmdlbLiK
x121Address: 198.176.123.101
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
dn: cn=Jesse Catu, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Jesse Catu
sn: Catu
telephonenumber: 1-812-855-7748
internationaliSDNNumber: 755-7748
postalcode: 1354
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
dn: cn=Joe Simms, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Joe Simms
sn: Simms
telephonenumber: 1-812-855-8395
internationaliSDNNumber: 755-8395
postalcode: 1514
telexnumber: 1-812-495-7860
seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company
dn: cn=Judy Simms, ou=In Flight Systems, ou=Home Town, o=Your Company
objectclass: top
objectclass: person
objectclass: organizationalPerson
cn: Judy Simms
sn: Simms
```

telephonenumber: 1-812-855-7352 postalcode: 1343 seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company dn: cn=Linda Carlesberg, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Linda Carlesberg sn: Carlesberg telephonenumber: 1-812-855-5492 internationaliSDNNumber: 755-5492 title: Purchasing Services seealso: cn=Cindy Jeffers, o=Your Company postalcode: 1109 facsimiletelephonenumber: 755-8985 dn: cn=Robert Dean, ou=In Flight Systems, ou=Home Town, o=Your Company objectclass: top objectclass: person objectclass: organizationalPerson cn: Robert Dean sn: Dean telephonenumber: 1-812-855-5703 internationaliSDNNumber: 755-5703 facsimiletelephonenumber: 1-812-855-5704 postalcode: 1701 facsimiletelephonenumber: 755-5704 seealso: cn=Maria Garcia, ou=In Flight Systems, ou=Home Town, o=Your Company

Appendix G. Supported server controls

The sections that follow describe the supported server controls.

IBMModifyDNTimelimitControl

- Name: IBMModifyDNTimelimitControl
- **Description:** Used by a client to request that a Modify DN operation be abandoned if the specified time limit for that operation has been exceeded.
 - Assigned object identifier: 1.3.18.0.2.10.10
- Target of control: Server
- Control criticality: Critical at client's option
- Values: The following ANSI.1 (Abstract Syntax Notation One) syntax describes the BER (Basic Encoding Rules) encoding of the control value.

```
ControlValue ::= SEQUENCE {
Time Limit INTEGER
}
```

For more information on ANS.1 and BER, go to the following Web site: ftp://ftp.rsa.com/pub/pkcs/ascii/layman.asc

Detailed description: This control is valid when sent on a client's Modify DN request. Modify DN operations may be long-running operations if they affect many entries in the directory (for example, if they rename an entry with a subtree containing many subordinate entries), so it may be desirable to limit the duration of the operation. The presence of this control on the Modify DN request causes the operation to be abandoned by the server if the number of seconds specified in the control value is exceeded. When the operation is abandoned, all changes to the directory associated with the Modify DN operation are undone.

IBMModifyDNRealignDNAttributesControl

- Name: IBMModifyDNRealignDNAttributesControl
- **Description:** Used by a client to request that a Modify DN operation be extended to realign attribute values for attributes of type DistinguishedName, and other specified attribute types known to contain distinguished names, with the new DN values established by the Modify DN operation for those DNs.
- Assigned object identifier: 1.3.18.0.2.10.11
- Target of control: Server
- Control criticality: Critical at client's option
- Values: There is no value; the *controlValue* field is absent.
- **Detailed description:** This control is valid when sent on a client's Modify DN request. Distinguished names which are renamed may be embedded in DN-syntax attributes throughout the directory contents. It may be desirable to replace the embedded values with their renamed counterparts (realignment). The presence of this control on the Modify DN request causes the server to realign matching attribute values in all DistinguishedName attributes, and all other attribute types whose attribute syntax is *Distinguished Name* (OID 1.3.6.1.4.1.1466.115.121.1.12), as well as in the attribute types of *aclEntry* and *entryOwner*, which are known to contain distinguished names. The server will evaluate whether the bound user has permission to modify the candidate attribute values, as determined by the appropriate access controls and the permissions granted by those access controls to the bound DN. If the permissions granted to the bound DN are sufficient to modify the candidate attribute values, those values will be realigned to match their respective new DN values. If any single access check fails, the entire operation fails, and all changes to the directory associated with the current Modify DN operation are undone.

manageDsalT

Name: manageDsalT

Description: Used on a request to suppress referral processing, thereby allowing the client to manipulate referral objects.

Assigned object identifier: 2.16.840.1.113730.3.4.2

· Target of control: Server · Control criticality: Critical

Values: There is no value; the controlValue field is absent.

• Detailed description: This control is valid when sent on a client's search, compare, add, delete, modify, or modify DN request. The presence of the control indicates that the server should not return referrals or search continuation references to the client. This allows the client to read or modify referral objects.

authenticateOnly

Name: authenticateOnly

Description: Used on an LDAP bind operation to indicate to the LDAP server that it should not attempt to find any group membership information for the client's bind DN.

Assigned object identifier: 1.3.18.0.2.10.2

Target of control: Server

Control criticality: Critical at client's option

Values: There is no value; the controlValue field is absent.

Detailed description: This control is valid when sent on an LDAP client's bind request to the LDAP server. The presence of this control on the bind request overrides alternate DN look-ups, extended group searching, and default group membership gathering, and causes the LDAP server to only authenticate the client's bind DN and not gather group information at all. This control is intended for a client who does not care about group memberships and subsequent complete authorization checking using groups, but is using the bind only for authentication to the LDAP server and fast bind processing.

IbmLDAPProxyControl

Name: IbmLDAPProxyControl

Description: Used to provide bind and connection information on extended operation requests that result in LDAP requests to any LDAP server. It is required on GetDnForUserid and GetPrivileges extended operation requests.

Assigned object identifier: 1.3.18.0.2.10.6 Target of control: EXOP backend of server

Control criticality: Critical

Values: The following ASN.1 (Abstract Syntax Notation One) syntax describes the BER (Basic

Encoding Rules) encoding of the control value.

```
ControlValue ::= SEQUENCE {
      BindInformation
                             [0] BindInfo OPTIONAL,
      ConnectInformation
                             [1] ConnectInfo OPTIONAL }
Where,
       ConnectInfo ::= LDAPURL
       BindInfo ::= SEQUENCE {
                               LDAPDN,
              Bind DN
                          AuthenticationChoice
              Auth
       AuthenticationChoice ::= CHOICE{
```

```
Simple
                       [0] OCTET STRING,
[3] SaslCredentials
}
Sas1Credentials ::= SEQUENCE {
                    LDAPString,
       Mechanism
                      OCTET STRING OPTIONAL
       Credentials
```

For more information on ASN.1 and BER, go to the following Web site:

ftp://ftp.rsa.com/pub/pkcs/ascii/layman.asc

Detailed description: This control provides information on extended operation requests that result in the use of the LDAP client to make LDAP requests to any LDAP server. The EXOP backend uses the connection information and the bind information specified in the control to establish an LDAP connection. Then, using the established connection, it issues additional LDAP requests to the server.

If the ConnectInformation is not specified, the EXOP backend attempts to open a connection to its local host using a default port of 389. Otherwise, it uses the LDAP URL specified to open a connection. The specified URL must have the following form:

ldap[s]://host[:port]

where:

- host is a DNS-style host name
- port is an optional port number
- Idaps causes the EXOP backend to open a secure LDAP connection. The LDAP server must be set up to use SSL and it cannot use the sslKeyRingPWStashFile option. See "Setting up for SSL/TLS" on page 43 for more information on SSL configuration.

If the BindInfo is not specified, the EXOP backend makes all of its LDAP requests anonymously. Otherwise, it uses the Bind DN and the AuthenticationChoice to bind to the LDAP server specified in the ConnectInformation. The EXOP backend does not support the SASL authentication choice which is described in the ASN.1.

Appendix H. Supported extended operations

The sections that follow describe the supported extended operations. For information on ASN.1 (Abstract Syntax Notation One) and BER (Basic Encoding Rules), go to the following Web site:

ftp://ftp.rsa.com/pub/pkcs/ascii/layman.asc

GetDnForUserid

- · Name: GetDnForUserid
- **Description:** Causes the EXOP backend to open a connection to an LDAP server with Policy Director data to retrieve all of a user ID's distinguished names.
- Assigned Object Identifier: 1.3.18.0.2.12.8
- Values: The following ASN.1 syntax describes the BER encoding of the request value.

Where,

Objectclass-name ::= OCTET STRING

Detailed description: Given a user ID and the required IbmLDAPProxyControl, the EXOP backend
opens a connection to the target LDAP server specified in the IbmLDAPProxyControl and retrieves all
of the specified user ID's distinguished names.

The search base in the request value establishes the sub-tree to search for the user ID's distinguished names. If this is not specified, the EXOP backend will perform a root DSE search to determine all of the naming contexts of the target LDAP server and proceed to search each naming context for the user ID's distinguished names. In addition to the search base, the user can specify optional object classes to filter distinguished names from the result.

- Response object identifier: 1.3.18.0.2.12.10
- Response description: Returned by EXOP backend when it receives a GetDnForUserid extended
 operations request.
- Response values: The following ASN.1 syntax describes the BER encoding of the response value.

ResponseValue ::= SEQUENCE OF Distinguished-name

where,

Distinguished-name LDAPDN

Response detailed description: When the EXOP backend receives a GetDnForUserid extended operation request and the required IbmLDAPProxyControl, it issues requests to the target LDAP sever specified in the IbmLDAPProxyControl to retrieve all of the specified user ID's distinguished names. The user may further filter the results by specifying a search base and object class names in the request value.

The following table summarizes some different error scenarios and the EXOP backend's response to such scenarios.

Error scenario	EXOP backend's response
Cannot find distinguished names	Returns an LDAP_NO_SUCH_OBJECT return code
Encounters an LDAP_NO_SUCH_OBJECT return code in its attempt to bind to the target LDAP server	Returns an LDAP_INAPPROPRIATE_AUTH return code

Error scenario	EXOP backend's response
· · · · · · · · · · · · · · · · · · ·	Returns these return codes encountered and a detailed message describing the point of failure

GetPrivileges

- Name: GetPrivileges
- Description: Causes the EXOP backend to open a connection to an LDAP server with Policy Director data and retrieve all of a subject's Policy Director privilege information.
- Assigned Object Identifier: 1.3.18.0.2.12.7
- Values: The following ASN.1 syntax describes the BER encoding of the request value.

```
RequestValue ::= SEQUENCE {
      Subject
                   LDAPDN,
      DomainName OCTET STRING OPTIONAL
```

- Detailed description: Given the subject and the required IbmLDAPProxyControl, the EXOP backend opens a connection to the target LDAP server specified in the IbmLDAPProxyControl and retrieves all of the specified subject's Policy Director data. If no domain name is specified, the EXOP backend assumes that the subject exists in the DEFAULT domain.
- Response object identifier: 1.3.18.0.2.12.9
- Response description: Returned by EXOP backend when it receives a GetPrivileges extended operations request.
- Response values: The following ASN.1 syntax describes the BER encoding of the response value.

```
ResponseValue ::= SEQUENCE {
                            OCTET STRING,
     DomainName
     SecLoginType
                            OCTET STRING,
                           OCTET STRING,
     PrincipalName
     SecPwdValid
                          BOOLEAN,
     SecAcctValid
                          BOOLEAN,
     UserUUID
                            SEQUENCE {
          Username
                                LDAPDN,
          UserUUID
                                OCTET STRING
     GroupInfo
                            SEQUENCE {
          NumberOfGroups
                               INTEGER,
                                SEQUENCE OF SEQUENCE {
          GroupUUIDs
                Groupname
                                   LDAPDN,
                GroupUUID
                                    OCTET STRING
     }
```

Response detailed description: When the EXOP backend receives a GetPrivileges extended operation request and the required IbmLDAPProxyControl, it issues requests to the target LDAP sever specified in the IbmLDAPProxyControl to retrieve all of the subject's Policy Director data. Refer to Policy Director Authorization Service for z/OS and OS/390: Customization and Use for descriptions of the fields returned in the response value.

The following table summarizes some different error scenarios and the EXOP backend's response to such scenarios.

Error scenario	EXOP backend's response
Cannot find any of the fields in the response value	Returns an LDAP_NO_SUCH_OBJECT return code
Retrieves more than one UserUUID or more than one GroupUUID per Groupname	Returns an LDAP_OTHER return code

Error scenario	EXOP backend's response
Encounters an LDAP_NO_SUCH_OBJECT return code in its attempt to bind to the target LDAP server	Returns an LDAP_INAPPROPRIATE_AUTH return code
Encounters any other unsuccessful return codes in its attempt to make LDAP requests to the target LDAP server	Returns these return codes encountered and a detailed message describing the point of failure

Start TLS

- Name: Start TLS Extended Request
- **Description:** Causes a non-secure connection to change to a secure connection.
- Assigned Object Identifier: 1.3.6.1.4.1.1466.20037
 - Values: None.
- Detailed description: The client may send the Start TLS extended request at any time after establishing an LDAP association, except in the following cases:
 - If a secure connection is already established, or
 - During a multi-stage SASL negotiation, or
 - If there are any outstanding LDAP operations on the connection.

The LDAP server will respond with an indication of whether the change to a secure connection is allowed. If accepted, the client is expected to immediately begin the secure protocol handshake.

The secure connection may be ended and a non-secure connection resumed by having the client cause a TLS closure alert to be sent to the server. Communication after receiving the TLS closure alert is over a non-secure connection. The client is considered to be in an anonymous authentication state.

- Response object identifier: 1.3.6.1.4.1.1466.20037
 - Response description: Upon receiving the Start TLS extended request, the server will return an extended response containing a response code indicating success or failure.
- Response values: For the successful response, no response value is returned. For an error response, a response message indicating the cause of the error is returned.
- · Response detailed description:

The following table summarizes some different error scenarios and the server's response to such scenarios.

Error scenario	Server response
Server accepts connection and can handle the request	Returns an LDAP_SUCCESS return code
A secure connection is already established	Returns an LDAP_OPERATIONS_ERROR return code
Secure connections are not supported by the server	Returns an LDAP_PROTOCOL_ERROR return code
There are outstanding operations on the connection	Returns an LDAP_PROTOCOL_ERROR return code
A multi-stage SASL negotiation is in progress	Returns an LDAP_PROTOCOL_ERROR return code

Appendix I. TDBM schema migration

This appendix describes the procedures for migrating the TDBM schema from an OS/390 V2R10, z/OS V1R1, z/OS V1R2, or z/OS V1R3 LDAP server to a z/OS V1R4 LDAP server.

Migrating TDBM schema from an OS/390 V2R10 or z/OS V1R1 LDAP server to a z/OS V1R4 LDAP server

- The following are the post-apply actions shipped with APAR OW53447 for customers who have OS/390 V2R10 or z/OS V1R1 LDAP server installed.
- This action is for customers who have configured the TDBM backend of the LDAP server.
- Customer action is required to make these updates to the schema in use by the TDBM backend in the LDAP server. These changes should be made by the LDAP Administrator.
- This PTF replaces the TDBM schema modifications found in OW50971. Do not apply the changes found in OW50971 related to updating the TDBM schema.
- This PTF updates the **schema-updates.ldif** schema file in **/usr/lpp/ldap/etc**. Copy the file from **/usr/lpp/ldap/etc** to **/etc/ldap**, update the *<suffix>* values in the file and apply the changes as directed below.

Step 1

- Instructions for updating schema associated with the TDBM backend:
- With this PTF, all customers are required to be at either the schema.user.ldif or schema.user.ldif plus
- schema.IBM.Idif level of schema modules loaded for TDBM. Previously, LDAP allowed users to load individual component schema modules (for example, RFC2256.Idif) or they could load the
- schema.user.ldif and schema.lBM.ldif files. Customers may now load either the schema.user.ldif file or
- both the **schema.user.Idif** and **schema.IBM.Idif** files, but after completing this HOLD action, should no longer apply the individual schema modules shipped with LDAP.
- If the **schema.user.Idif** file was used to create the schema in use by TDBM and the **schema.IBM.Idif** file will not be applied, then continue with Step 2 below.
- If the **schema.user.Idif** file was used to create the schema in use by TDBM and you intend to spply the **schema.IBM.Idif** file, then apply the updates in Step 2 for the **schema.user.Idif** file and then apply the **schema.IBM.Idif** file which was shipped with OW50971.
- If the **schema.user.ldif** and **schema.IBM.ldif** files were used to create the schema in use by TDBM, then also add the **CommServer.ldif** schema file. Then continue with Step 2 below.
- If any individual schema modules were applied rather than the **schema.user.ldif** or **schema.user.ldif** plus **schema.IBM.ldif** schema modules, then load any remaining schema modules to get to either the
- schema.user.ldif or schema.user.ldif plus schema.lBM.ldif level. Use the individual schema modules
- which were shipped with OW50971. Attempting to reload any schema modules which have previously
- been applied will result in a 'type or value exists' message. This generally indicates that this module has previously been applied and the message may be ignored.
- The following list shows the order in which the OS/390 V2R10 or z/OS V1R1 schema modules should be loaded to get to the **schema.user.ldif** or **schema.user.ldif** plus **schema.lBM.ldif** levels.
- Use the **Idapmodify** command (from the USS shell):

	<pre>ldapmodify -h <host> -p <port> -D <admindn> -w <adminpw> -f /etc/ldap/<filename.ldif></filename.ldif></adminpw></admindn></port></host></pre>
	specifying the name of the schema module to be loaded from the list below.
	If the schema module has been applied to the TDBM schema, then it does not need to be applied again. If it is applied again, there may be messages stating that 'type or value exists'. These messages may be ignored.
	See Chapter 16, "LDAP directory schema for TDBM" on page 159 for instructions on how to apply modifications to the TDBM schema.
	Components of schema.user.ldif The following files have no pre-requisites and can be loaded in any order: ChangeLog.ldif ComponentBroker.ldif nisSchema.ldif RFC2252.ldif RFC2256.ldif
	The following file requires RFC2256.ldif to be loaded before it can be loaded: • Netscape.ldif
	The following file requires RFC2256.ldif, Netscape.ldif, and OtherStandard.ldif to be loaded before it can be loaded: • X.520.ldif
	The following files require X.520.ldif to be loaded before they can be loaded: • Netscape-V2.ldif • System.ldif • UniversalMessaging.ldif
	The following file requires System.Idif to be loaded before it can be loaded: • System-V2.Idif
	At this point, the schema loaded is equivalent to schema.user.ldif . If schema.IBM.ldif is to be added to the schema, add just that module to the schema entry by using the ldapmodify command and specifying the /etc/ldap/schema.IBM.ldif file. Then, proceed to Step 2.
	If individual components of schema.IBM.Idif were loaded, then determine which file(s) have been loaded and proceed to load the remaining files in the order as directed below.
	Components of schema.IBM.Idif The following file has no pre-requisites and can be loaded at any time: • CommServer.Idif
	The following file requires the equivalent of schema.user.ldif to be loaded before it can be loaded: • DMTF.ldif
	The following files require DMTF.Idif to be loaded before they can be loaded: • MCI.Idif • IBM.Idif

	The following files require IBM.Idif to be loaded • DB2.Idif	before they can be loaded:
i	Kerberos-V1.Idif	
i	MetaDirectory.ldif	
İ	• NFI.Idif	
	OnDemandServer.Idif	
	• RACF.Idif	
	RegisteredSoftware.ldif	
	The following file requires OnDemandServer.Idi • UNIX.Idif	f to be loaded before it can be loaded:
	Step 2	
	this PTF should be applied to the schema. Some	chema.user.ldif or schema.user.ldif plus gh 16 found in the schema-updates.ldif file shipped with e of these updates are manual tasks. The comments refully reviewed to determine what updates need to be
	The following table shows which updates relate t	to which previously shipped APAR:
	APAR OW47125	Updates 1 through 6
	APAR OW48437	Updates 7 through 9
	APAR OW46344 APAR OW47596	Update 10 Update 11
i	APAR OW53447 (this)	Updates 12 through 15
İ	APAR OW50024	Update 16
	·	ously applied and the schema updated as described in the ciated with those APARS do not need to be applied again.
	If the schema module is applied again, there ma messages may be ignored.	y be messages stating 'type or value exists'. These
	See the schema-updates.ldif file for specific information concerning each update.	
	Migrating TDBM schema from a to a z/OS V1R4 LDAP server	z/OS V1R2 or z/OS V1R3 LDAP server
	The following are the post-apply actions shipped V1R2 or z/OS V1R3 LDAP server installed.	with APAR OW53447 for customers who have a z/OS
	This action is for customers who have configured	d the TDBM backend of the LDAP Server.
	Customer action is required to make these updar LDAP Server. These changes should be made b	tes to the schema in use by the TDBM backend in the y the LDAP Administrator.
	This PTF replaces the TDBM schema modification	ons found in OW50971. Do not apply the changes found in

OW50971 related to updating the TDBM schema.

	This PTF updates the schema-updates.ldif schema file in /usr/lpp/ldap/etc . Copy the file from /usr/lpp/ldap/etc to /etc/ldap , update the <suffix></suffix> values in the file, and apply the changes as directed below.
	Instructions for updating schema associated with the TDBM backend.
	If z/OS V1R2 or z/OS V1R3 LDAP schema files were applied to create the TDBM schema, then perform the applicable tasks in Step 1 using the z/OS V1R2 or z/OS V1R3 schema files. Step 2 does not need to be performed.
	If the LDAP server has been migrated from OS/390 V2R10 or z/OS V1R1, perform the tasks in Step 1 that apply using the z/OS V1R2 or z/OS V1R3 schema files and then proceed to Step 2.
	Step 1 With this PTF, all customers are required to be at either the schema.user.ldif or schema.user.ldif plus schema.lBM.ldif level of schema modules loaded for TDBM. Previously, LDAP allowed users to load individual component schema modules (for example, RFC2256.ldif) or they could load the schema.user.ldif and schema.lBM.ldif files. Customers may now load either the schema.user.ldif file or both the schema.user.ldif and schema.lBM.ldif files, but after completing this HOLD action, should no longer apply the individual schema modules shipped with LDAP.
	If the OS/390 V2R10 or z/OS V1R1 schema.user.ldif file was used to create the schema in use by TDBM and the schema.lBM.ldif file will not be applied, then continue with Step 2 below.
	If the OS/390 V2R10 or z/OS V1R1 schema.user.ldif file was used to create the schema in use by TDBM and you intend to apply the schema.lBM.ldif file, then apply the updates in Step 2 for the schema.user.ldif file and then apply the schema.lBM.ldif file which was shipped with OW50971.
	If the OS/390 V2R10 or z/OS V1R1 schema.user.ldif and schema.IBM.ldif files were used to create the schema in use by TDBM, then also add the CommServer.ldif schema file. Then continue with Step 2 below.
	If the z/OS V1R2 or z/OS V1R3 schema.user.ldif file was used to create the schema in use by TDBM and the schema.lBM.ldif file will not be applied, then there are no additional updates at this time.
	If the z/OS V1R2 or z/OS V1R3 schema.user.ldif file was used to create the schema in use by TDBM and you intend to apply the schema.lBM.ldif file, then apply the schema.lBM.ldif file which was shipped with z/OS V1R2 or z/OS V1R3. No additional updates will need to be applied.
	If the z/OS V1R2 or z/OS V1R3 schema.user.ldif and schema.lBM.ldif files were used to create the schema in use by TDBM, then there are no additional updates at this time.
	If any individual schema modules were applied rather than the schema.user.ldif or schema.user.ldif plus schema.lBM.ldif schema modules, then load any remaining schema modules to get to either the schema.user.ldif or schema.user.ldif plus schema.lBM.ldif level. Use the individual schema modules which were shipped with z/OS V1R2 or z/OS V1R3 LDAP. Attempting to reload any schema modules which have previously been applied will result in a 'type or value exists' message. This generally indicates that this module has previously been applied and the message may be ignored.
	The following list shows the order in which the z/OS V1R2 or z/OS V1R3 schema modules should be loaded to get to the schema.user.ldif or schema.user.ldif plus schema.lBM.ldif levels.
	Use the Idapmodify command (from the USS shell): ldapmodify -h <host> -p <port> -D <admindn> -w <adminpw> -f /etc/ldap/<filename.ldif></filename.ldif></adminpw></admindn></port></host>

	specifying the name of the schema module to be loaded from the list below.
	If the schema module has been applied to the TDBM schema, then it does not need to be applied again. If it is applied again, there may be messages stating that 'type or value exists'. These messages may be ignored.
	See Chapter 16, "LDAP directory schema for TDBM" on page 159 for instructions on how to apply modifications to the TDBM schema.
	Components of schema.user.ldif The following files have no pre-requisites and can be loaded in any order: ChangeLog.ldif ComponentBroker.ldif MS.ActiveDirectory.ldif nisSchema.ldif OtherStandard.ldif RFC2252.ldif RFC2256.ldif
	The following files require RFC2256.ldif to be loaded before they can be loaded: • Netscape.ldif • RFC2587.ldif • RFC2713.ldif • RFC2714.ldif
	The following file requires MS.ActiveDirectory.ldif to be loaded before it can be loaded: • SecurityIdentities.ldif
	The following file requires RFC2256.ldif, Netscape.ldif, and OtherStandard.ldif to be loaded before it can be loaded: • System.ldif
	The following file requires System.Idif to be loaded before it can be loaded: • X.520.Idif
	The following files require X.520.ldif to be loaded before they can be loaded: • Netscape-V2.ldif • UniversalMessaging.ldif • EntrustPKIV4.ldif
	The following file requires EntrustPKIV4.Idif to be loaded before it can be loaded: • EntrustPKIV5.Idif
	The following file requires System.Idif to be loaded before it can be loaded: • System-V2.Idif
	At this point, the schema loaded is equivalent to schema.user.ldif . If schema.lBM.ldif is to be added to the schema, add just that module to the schema entry by using the Idapmodify command and specifying the /etc/Idap/schema.IBM.ldif file.

	If individual components of schema.IBM.Idif were loaded, then determine which file(s) have be and proceed to load the remaining files in the order as directed below.	en loaded	
	Components of schema.IBM.Idif The following file has no pre-requisites and can be loaded at any time: CommServer.Idif NativeAuthentication.Idif		
	The following files require the equivalent of schema.user.ldif to be loaded before they can be • DMTF.ldif • WebSphereNaming.ldif	loaded:	
	The following files require DMTF.Idif to be loaded before they can be loaded: • MCI.Idif • MCI.Idif • IBM.Idif		
	The following files require IBM.Idif to be loaded before they can be loaded: DB2.Idif Kerberos-V1.Idif MetaDirectory.Idif NFI.Idif OnDemandServer.Idif PolicyDirector.Idif RACF.Idif RegisteredSoftware.Idif		
 	The following files require OnDemandServer.ldif to be loaded before they can be loaded: • ManagedSystemInfrastructure.ldif • UNIX.ldif		
 	The following file requires RACF.Idif to be loaded before it can be loaded: • RACF.2.Idif		
	Step 2 If z/OS V1R2 or z/OS V1R3 LDAP schema files were applied to create the TDBM schema, and the applicable tasks in Step 1 using the z/OS V1R2 or z/OS V1R3 schema files have been performed, then Step 2 does not need to be performed.		
	If the LDAP server has been migrated from OS/390 V2R10 or z/OS V1R1, proceed with Step 2.		
	After updating the loaded schema to either the schema.user.ldif or schema.user.ldif plus schema.IBM.ldif level, schema updates 1 through 16 found in the schema-updates.ldif file shipped with this PTF should be applied to the schema. Some of these updates are manual tasks. The comments contained in schema-updates.ldif should be carefully reviewed to determine what updates need to be made to the schema.		
	The following table shows which updates relate to which previously shipped APAR:		
	APAR OW47125 Updates 1 through 6 APAR OW48437 Updates 7 through 9		

APAR OW46344 Update 10 Update 11 APAR OW47596 Updates 12 through 15 APAR OW53447 (this) Update 16 APAR OW50024

| If one or more of these APARS have been previously applied and the schema updated as described in the associated HOLD action, then the updates associated with those APARS do not need to be applied again.

If the schema module is applied again, there may be messages stating 'type or value exists'. These messages may be ignored.

| See the **schema-updates.ldif** file for specific information concerning each update.

Appendix J. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS^{TM} enable users to:

- · Use assistive technologies such as screen-readers and screen magnifier software
- · Operate specific or equivalent features using only the keyboard
- · Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen-readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to z/OS: TSO/E Primer, z/OS: TSO/E User's Guide, and z/OS: ISPF User's Guide Volume I for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing 2-31 Roppongi 3-chome, Minato-ku Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact: but do not indicate that it is the legal department.

IBM Corporation Mail Station P300 2455 South Road Poughkeepsie, NY 12601-5400 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© International Business Machines Corporation 1999. Portions of this code are derived from IBM Corp. Sample Programs. Copyright IBM Corp. 1999. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	DB2®	OS/400 [®]	RMF™
AIX®	DRDA®	Parallel Sysplex®	SecureWay [®]
CICS®	OpenEdition®	RACF®	WebSphere®
DATABASE 2	OS/390 [®]	RETAIN®	z/OS

Lotus® is a trademark of Lotus Development Corporation in the United States, other countries, or both.

 $Microsoft^{@}$, $Windows^{@}$, $Windows\ NT^{@}$, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Kerberos is a trademark of the Massachusetts Institute of Technology (MIT).

 $\mathsf{UNIX}^{\scriptscriptstyle{(\! g)}}$ is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be the trademarks or service marks of others.

Bibliography

This bibliography provides a list of publications that are useful when implementing the LDAP server product.

IBM z/OS SecureWay Security Server publications

- z/OS: Security Server LDAP Client Programming, SC24-5924
- z/OS: RACF Security Administrator's Guide, SA22-7683
- z/OS: Security Server RACF Migration, GA22-7690
- z/OS: Security Server RACF System Programmer's Guide, SA22-7681
- z/OS: Security Server RACF Command Language Reference, SA22-7687
- z/OS: Security Server RACF Callable Services, SA22-7691
- z/OS: Security Server Network Authentication Service Administration, SC24-5926
- z/OS: Security Server Network Authentication Service Programming, SC24-5927

IBM C/C++ language publications

- z/OS: C/C++ Programming Guide, SC09-4765
- z/OS: C/C++ Run-Time Library Reference, SA22-7821

IBM DB2 publications

- DB2 for OS/390 and z/OS: ODBC Guide and Reference, GC26-9941
- DB2 for OS/390 and z/OS: Application Programming and SQL Guide, SC26-9933
- DB2 for OS/390 and z/OS: Installation Guide, GC26-9936
- DB2 for OS/390 and z/OS: Command Reference, SC26-9934
- DB2 for OS/390 and z/OS: Messages and Codes, GC26-9940
- DB2 for OS/390 and z/OS: SQL Reference, SC26-9944
- DB2 for OS/390 and z/OS: Data sharing: Planning and Administration, SC26-9935
- DB2 for OS/390 and z/OS: Utility Guide and Reference, SC26-9945

 DB2 for OS/390 and z/OS: ODBC Guide and Reference, GC26-9941

IBM z/OS Cryptographic Service publications

- z/OS: Open Cryptographic Services Facility Application Programming, SC24-5899
- z/OS: ICSF Administrator's Guide, SA22-7521

Other IBM publications

- z/OS: Communications Server: IP Configuration Guide, SC31-8775
- z/OS: Communications Server: IP Configuration Reference, SC31-8776
- z/OS: Program Directory, GI10-0670
- z/OS: System Secure Sockets Layer Programming, SC24-5901
- z/OS: UNIX System Services Planning, GA22-7800
- z/OS: Parallel Sysplex Overview: An Introduction to Data Sharing and Parallelism, SA22-7661
- z/OS: MVS Setting Up a Sysplex, SA22-7625
- z/OS: SDSF Operation and Customization, SA22-7670
- z/OS: DCE Command Reference, SC24-5909
- z/OS: Planning for Installation, GA22-7504
- z/OS: Introduction and Release Guide, GA22-7502
- z/OS: DCE Application Development Guide: Directory Services, SC24-5906
- z/OS: Licensed Program Specifications, GA22-7503
- z/OS: Collection, SK3T-4269
- z/OS: Information Roadmap, SA22-7500
- Policy Director Authorization Service for z/OS and OS/390: Customization and Use, SC24-6040
- ServerPac: Installing Your Order

© Copyright IBM Corp. 1998, 2002 **453**

Glossary

This glossary defines technical terms and abbreviations used in z/OS LDAP documentation. If you do not find the term you are looking for, refer to the index of the appropriate z/OS LDAP manual or view *IBM Dictionary of Computing*, available from

www.ibm.com/ibm/terminology

This glossary includes terms and definitions from:

- IBM Dictionary of Computing, SC20-1699.
- Information Technology—Portable Operating System Interface (POSIX), from the POSIX series of standards for applications and user interfaces to open systems, copyrighted by the Institute of Electrical and Electronics Engineers (IEEE).
- American National Standard Dictionary for Information Systems, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.
- Information Technology Vocabulary, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1.SC1).
- CCITT Sixth Plenary Assembly Orange Book, Terms and Definitions and working documents published by the International Telecommunication Union, Geneva, 1978.
- · Open Software Foundation (OSF).

Α

access control. Ensuring that the resources of a computer system can be accessed only by authorized users in authorized ways.

access control list (ACL). Data that controls access to a protected object. An ACL specifies the privilege attributes needed to access the object and the permissions that may be granted, to the protected object, to principals that possess such privilege attributes.

ACL. Access control list.

attribute. Information of a particular type concerning an object and appearing in an entry that describes the

object in the directory information base (DIB). It denotes the attribute's type and a sequence of one or more attribute values, each accompanied by an integer denoting the value's syntax.

B

backend. A subsystem of the LDAP server which implements access to a persistent storage mechanism for information.

C

certificate. Used to prove your identity. A secure server must have a certificate and a public-private key pair. A certificate is issued and signed by a Certificate Authority (CA).

cipher. A method of transforming text in order to conceal its meaning.

CKDS. Cryptographic Key Data Set.

client. A computer or process that accesses the data, services, or resources of another computer or process on the network. Contrast with *server*.

configuration. The manner in which the hardware and software of an information processing system are organized and interconnected.

Cryptographic Key Data Set (CKDS). (1) A data set that contains the encrypting keys used by an installation. (2) In ICSF, a VSAM data set that contains all the cryptographic keys. Besides the encrypted key value, an entry in the cryptographic key data set contains information about the key.

cryptography. (1) The transformation of data to conceal its meaning. (2) In computer security, the principles, means, and methods for encrypting plaintext and decrypting ciphertext. (3) In ICSF, the use of cryptography is extended to include the generation and verification of MACs, the generation of MDCs and other one-way hashes, the generation and verification of PINs, and the generation and verification of digital signatures.

D

daemon. A long-lived process that runs unattended to perform continuous or periodic system-wide functions such as network control. Some daemons are triggered automatically to perform their task; others operate periodically.

Data Encryption Standard (DES). In computer security, the National Institute of Standards and Technology (NIST) Data Encryption Standard, adopted by the U.S. government as Federal Information Processing Standard (FIPS) Publication 46, which allows only hardware implementations of the data encryption algorithm.

data hierarchy. A data structure consisting of sets and subsets such that every subset of a set is of lower rank than the data of the set.

data model. (1) A logical view of the organization of data in a database. (2) In a database, the user's logical view of the data in contrast to the physically stored data, or storage structure. (3) A description of the organization of data in a manner that reflects information structure of an enterprise.

database. A collection of data with a given structure for accepting, storing, and providing, on demand, data for multiple users.

Database 2 (DB2). An IBM relational database management system.

DB2. Database 2.

DES. Data Encryption Standard (DES).

directory. (1) A logical unit for storing entries under one name (the directory name) in a CDS namespace. Each physical instance of a directory is called a replica. (2) A collection of open systems that cooperates to hold a logical database of information about a set of objects in the real world.

directory schema. The set of rules and constraints concerning directory information tree (DIT) structure, object class definitions, attribute types, and syntaxes that characterize the directory information base (DIB).

directory service. The directory service is a central repository for information about resources in a distributed system.

distinguished name (DN). One of the names of an object, formed from the sequence of RDNs of its object entry and each of its superior entries.

DN. Distinguished name.

F

environment variable. A variable included in the current software environment that is available to any called program that requests it.

ICSF. Integrated Cryptographic Service Facility.

Integrated Cryptographic Service Facility (ICSF). A licensed program that runs under z/OS and provides access to the hardware cryptographic feature for programming applications. The combination of the hardware cryptographic feature and ICSF provides secure high-speed cryptographic services

J

JCL. Job control language.

Job control language (JCL). A control language used to identify a job to an operating system and to describe the job's requirements.

K

key generator utility program (KGUP). A program that processes control statements for generating and maintaining keys in the cryptographic key data set.

KGUP. Key generator utility program.

LDAP. Lightweight Directory Access Protocol.

Lightweight Directory Access Protocol (LDAP). A client/server protocol for accessing a directory service.

M

master replica. The first instance of a specific directory in the namespace. After copies of the directory have been made, a different replica can be designated as the master, but only one master replica of a directory can exist at a time.

MD5. Message Digest 5. A hash algorithm.

MKKF. Make Key File.

MKKF utility. A command-line utility used to create public/private key pairs and certificate requests, receive certificate requests into a key ring, and manage keys in a key ring.

0

object class. An identified family of objects that share certain characteristics. An object class can be specific to one application or shared among a group of applications. An application interprets and uses an entry's class-specific attributes based on the class of the object that the entry describes.

OCSF. Open Cryptographic Services Facility.

ODBC. Open database connectivity.

Open Cryptographic Services Facility (OCSF). A derivative of the IBM Keyworks technology which is an implementation of the Common Data Security Architecture (CDSA) for applications running in the UNIX System Services environment.

z/OS Cryptographic Services. A z/OS offering that supplies a set of interfaces for cryptographic functions.

private key. Used for the encryption of data. A secure server keeps its private key secret. A secure server sends clients its public key so they can encrypt data to the server. The server then decrypts the data with its private key.

public key. Used for the encryption of data. A secure server makes its public key widely available so that its clients can encrypt data to send to the server. The server then decrypts the data with its private key.

R

RACF. Resource Access Control Facility.

RDN. Relative distinguished name.

referral. An outcome that can be returned by a directory system agent that cannot perform an operation itself. The referral identifies one or more other directory system agents more able to perform the operation.

relative distinguished name (RDN). A component of a DN. It identifies an entry distinctly from any other entries which have the same parent.

replica. A directory in the CDS namespace. The first instance of a directory in the namespace is the master replica. See master replica.

replication. The making of a shadow of a database to be used by another node. Replication can improve availability and load-sharing.

Resource Access Control Facility (RACF). An IBM licensed program, that provides for access control by identifying and verifying the users to the system, authorizing access to protected resources, and logging the detected unauthorized access to protected resources.

S

SASL. Simple Authentication Security Layer.

schema. See directory schema.

Secure Sockets Layer (SSL) security. A facility used to protect LDAP access.

server. On a network, the computer that contains programs, data, or provides the facilities that other computers on the network can access. Contrast with client.

SHA. Secure Hash Algorithm. A hash algorithm required for use with the Digital Signature Standard.

Simple Authentication Security Layer (SASL). Refers to a method of binding using authentication information outside the client and server.

SLAPD. A stand-alone LDAP daemon.

SPUFI. SQL Processor Using File Input.

SQL Processor Using File Input (SPUFI). A facility of the TSO attachment subcomponent that enables the DB2I user to run SQL statements without embedding them in an application program.

SQL. Structured Query Language.

SSL. Secure Sockets Layer.

Structured Query Language (SQL). A standardized language for defining and manipulating data in a relational database.

thread. A single sequential flow of control within a process.

Time Sharing Option (TSO). An operating system option that provides interactive time sharing from remote terminals.

Transport Layer Security (TLS). A security protocol that provides communication privacy over the Internet. The protocol alllows client/server applications to communicate in a way that is designed to prevent eavesdropping, tampering, or message forgery. TLS is based upon SSL Version 3.0

TSO. Time Sharing Option.

U

UCS Transformation Format (UTF). The LDAP Version 3 protocol specifies that data is passed between client and server in the UTF-8 character set.

UTF. UCS Transformation Format.



X.500. The CCITT/ISO standard for the open systems interconnection (OSI) application-layer directory. It allows users to register, store, search, and retrieve information about any objects or resources in a network or distributed system.

Index

Special characters	Access Control List (ACL) (continued) information, retrieving 258
_passwd ()	modifying a group for 268
errno values returned by 205	modifying in TDBM 261
/etc directory	modifying owner for entry 264
using 35	override example 257
/usr/lpp/ldap directory	ownerPropagate attribute 254
LDAP server install directory 35	ownerSource attribute 254
" (quotation marks) 52 ' (apostrophe) 156	propagation 254, 256
> (greater than sign) 156	requested attributes 256
# (pound sign support in SDBM) 211	searching 256
# (pound sign) 156	access protection 70, 73
+ (plus sign) 156	accessibility 447
= (equal sign) 156	ACL (Access Control List)
(-1	See Access Control List (ACL) 6
	aclEntry attribute
Numerics	description 250
7-bit ASCII 54, 61	migrating 106
7 51.71.5011	aclPropagate attribute
	description 253 aclSource attribute
A	
abandon behavior 298	description 253 activating
ABSTRACT object class type 165, 172	ICSF for password encryption 16
access	activity log
determining 254	modify DN 203
access classes	activity logging 100
attribute 251	adminDN option 53, 54, 80
determining 250	administration
permissions 252	restricting access 6
specifying for TDBM 165, 171	administrator
access control	configuring DN of 80
See also Access Control List (ACL)	password, specifying 55
attributes 249	roles for Idapcnf 22
configuring for Kerberos 227	specifying DN of 54
groups 258	adminPW option 53, 55, 80
LDAP server capability 7	advanced configuration options 27
using 249	alternate server
using RACF 205	specifying 55
access control and ownership	altServer option 55
modify DN 189	analyzing
Access Control List (ACL)	schema errors 174
aclEntry attribute 250	anonymous searches 256
aclPropagate attribute 253 aclSource attribute 253	APF authorization
attribute classes 256	for LDAP server 37 for SSL 37
creating a group for 266	Idapenf utility 21
creating a group for 250 creating in TDBM 259	migrating 108
creating owner for entry 263	migrating 100
deleting a group for 268	RDBM backend 37
deleting in TDBM 262	TDBM backend 37
deleting owner for entry 266	apostrophe 156
description 6, 249	arguments
determining access from 254	configuration 52
entryOwner attribute 253	arranging information 4
examples 257	ASCII, 7-bit 54, 61
filters 256	attribute
groups 258	LDAP server user ID 35

attribute (continued)	authorization
object class 5	native authentication 233
syntaxes 166	authorization, APF
attribute classes	See APF authorization 37
searching 256	AUXILIARY object class type 165, 172
attribute option 55	7.67.12.1
attribute types	
description 161	В
format 170	
usage 170	backend
	See also EXOP backend
attribute value changes 106	See also RDBM backend
attribute,ibm-entryuuid 9	See also SDBM backend
attributes	See also TDBM backend
access allowed for 250	database definitions 51
access classes 250, 251	databases 7
access control 249	DB2, using 13
aclEntry 250	multiple 7
aclPropagate 253	options for 53
aclSource 253	preparing for TDBM, RDBM, SDBM 33
cn 271	referral objects 278
deleting in SDBM 222	backend utilities, DB2 125, 126
description 271	backing store, DB2 7
determining 155	backslash character
entryOwner 253	DN syntax 156
jpegPhoto 4	using in configuration file 52
Kerberos 226	backup of master server 270
LDAP schema 167	benefits of replication 270
mail 4	bibliography 453
mandatory for replica object 270	bind, SASL 8
multi-valued ref 277	binding
multi-valued with RACF 216	6
normalizing 298	authenticateOnly server control 432 CLI plan 14
optional for replica object 271	in TDBM using native authentication 233
ownerpropagate 254	Kerberos 225
ownerSource 254	
ref 277	blank spaces
replicaBindDN 270	using in configuration file 52
replicaBindMethod 271	using in DNs 156
·	book organization xvii
replicaCredentials 270	books, related 453
replicaHost 270	buffer pools
replicaPort 271	setting up for 14
replicaUpdateTimeInterval 271	building
replicaUseSSL 271	directory namespace 288
requested 256	
returning lowercase 298	
searching 256	C
seeAlso 271	CAF (Call Attachment Facility)) 15
syntaxes 155	Call Level Interface (CLI)
type 160	plan, binding 14
type files, converting 177	setting up for 13
type, defining 55	capabilities of LDAP server 6
attributeTypes schema attribute 164, 171	capturing performance information 100
attrOverflowSize option 55	certificate
authenticateOnly server control 57, 432	authenticating 68
authentication	
client 8	client 8
Kerberos 225	digital
server 8, 46, 68	server 46
authentication bind	changed information xxiii
CRAM-MD5 and DIGEST-MD5 8	changed reason codes 298
C. C MDO GIIG DIOLOT MIDO	changes to book xxii, xxiii
	changes to interfaces 123

checklist for configuration file 52 CICS (Customer Information Control System) updating related attributes 216 ciphers specifications 68	configuring (continued) considerations 74 EXOP backend 32, 84 HCD 11 ICSF for password encryption 16
classes, access attribute 251 determining 250 permissions 252 specifying for TDBM 165, 171	LDAP for msys 29 LDAP server capability 7 OCSF for password encryption 16 operating in multi-server mode 77 operating in single-server mode 76
clear text password 47, 65, 148 CLI (Call Level Interface)	password encryption 82 RDBM backend 32
See Call Level Interface (CLI) 13	replica server 273
CLI initialization file name 33 client, LDAP	roadmap for 31 running with SDBM 42
See LDAP client 297	scenarios 82
code page IBM-1047 51	SDBM backend 31, 83
command-line parameters	SLAPD 75
for LDAP server 94	SSL 32, 82
commands	TDBM backend 32, 82, 83, 85
cp 39	using Idapenf utility 17
comment lines 52	WebSphere Application Server 11
commThreads option 55, 73, 107	connection
communication thread pool 55	group 215
complex modify DN replication 269	connections
Component Broker	specifying number of 62
configuring for 11 concurrent database instances 7	specifying timeout 57 control, program
configuration	See program control 37
specifying value for filename 53	controls, server
configuration file	See server controls 9
administrator DN, specifying 80	conventions used in book xvii
alternate 75	converting
arguments 52	See migrating 103
copying 39	copying
creating 51	configuration files 39
data set 126	files into data sets 97
data set versions of 97	cp command 39
default referral 278	CRAM-MD5 and DIGEST-MD5 243
deprecated options 73 format 51	configuration parameter 244 TDBM backend 243
Kerberos authentication 225	CRAM-MD5 and DIGEST-MD5 authentication 8
locating 51	creating
master server 275	ACL 259
options 53	DB2 CLI initialization file 14
options checklist 52	group for ACL 266
password encryption 47	owner for entry 263
password, specifying 80	referral entries 277
scenarios 82	slapd.conf 51
setting SSL keywords 274	SYSTEM file 131
slapd.conf 51, 371	user ID for LDAP server 35
specifying as data set name 96 specifying as DD name 96	versions of envvars file 97 critical access class 251
using to configure 75	crypt encryption format 48
configuration option	crypt encryption method 64
removing verifySchema 110	orypt choryption motiled or
configuration utility	
See Idapcnf utility 17	D
configuring	DAP (Directory Access Protocol)
access control, Kerberos 227	See Directory Access Protocol (DAP) 6
Component Broker 11	data key 65

data model		defining (continued)
LDAP 155		new TDBM schema elements 172
data set		participation in native authentication 233
accepting files as 97		started task 93
allocating for Idif2tdbm 129		terms 455
CLI Initialization sequential, specifying	57	definitions of terms 455
specifying configuration file as 96		deleting
specifying for configuration file 126		access control group 258
DATABASE 2 (DB2)		ACL 262
backend utilities, running 125, 126		attributes in SDBM 222
backends 7		group for ACL 268
backing store 7, 97		owner for entry 266
CLI initialization file 14		deprecated configuration options 73
creating file for 97		DES encryption format 49
creating TDBM database 40		DES encryption method 64, 270
granting resource authorizations 41		digestRealm option 56
input file for Idapcnf 27		digital certificate 46
installing 13		directory
referral objects 278		description 287
running 13		hierarchy example 4
server location, specifying 67		identifying entry in 155
tables, partitioning for TDBM 41		namespace 287
database administrator		of slapd.conf file 51
role for Idapcnf 23		planning content 11
database name 33		schema, TDBM 159
database option 56		updating 6
database owner 33		Directory Access Protocol (DAP)
databasename option 56		defining 6
databases		directory namespace
backend 7		example of building 288
multiple instances 7		organizing 287
DB2 (DATABASE 2)		directory schema
See DATABASE 2 (DB2) 13		migrating RDBM to TDBM 176
DB2 CLI		updating for Kerberos 226
See Call Level Interface (CLI) 13		directory service
DB2 subsystem ID 33		description 3
DB2I (DB2 Interactive) 40 db2Idif utility		directory, default LDAP 35 disability 447
encrypting user passwords 48		displaying
running from JCL 125		schema entry 175, 176
running in shell 125		distinguished name (DN)
running in TSO 126		administrator 54, 80
using with replica server 274		description 4, 155
db2pwden utility		group 258
description 148		length, maximum 155
encrypting passwords 48		master server 61
running from JCL 147		pound sign (#) 211
running in z/OS shell 147		RACF-style 156
dbuserid option 56		ref attribute 277
debug levels 94, 98		referencing by 5
debugging facility		reflecting 253, 254
turning on and off 98		retrieving with EXOP backend 247, 435
default		syntax 156
environment variable file 97		UTF-8 characters in 54
LDAP directory 35		DLL (dynamic link libraries)
mapping, Kerberos 227		See dynamic link libraries (DLL) 93
referral 65, 278		DN (distinguished name)
defining		See distinguished name (DN) 4
default referral 278		DN modify
Kerberos identity 36		ownership changes 193
LDAP server user ID 35		-

documentation	environment variables (continued)
migration 103	NLSPATH 151
related 453	setting for utilities 125
domain component naming 156	environments
DSE, root	operating, Kerberos 231
See root DSE 9	envvars file
DSNAOINI environment variable 57	data set versions of 97
DSNAOINI file 14, 97	file, default 97
dsnaoini option 57	LDAP_SLAPD_ENVVARS_FILE 97
dynamic debugging	PATH 97
using 98	setting 97
dynamic link libraries (DLL)	equal sign 156
program controlling for OCSF 37	errno values returned by _passwd() 205
using in startup 93	error codes 299
dynamic schema	errors
See also schema (TDBM)	schema, analyzing 174
TDBM 159	escape mechanism for UTF-8 characters 54
using TDBM backend 8 dynamic workload management	etc directory using 35
configuring 77	example
configuring 77 configuring LDAP server for 7	CRAM-MD5 and DIGEST-MD5 244
multi-server mode with 76	SDBM schema 220
main-server mode with 70	examples
	ACL 257
E	aclEntry attribute 256
editing LDIF files 291	adding a group to RACF 221
elements	adding user to RACF 220
schema, defining new 172	attribute definition 251
enabling	building directory namespace 288
SSL 46	configuration file 289, 294
encryption	configuration scenarios 82
for communication channel 46	configuring EXOP 84
encryption format 49	configuring multiple servers 85
encryption, password	configuring SDBM and TDBM 83
See password encryption 64	configuring TDBM, SSL, and password
entries	encryption 82
access allowed for 250	connecting user to group in RACF 221
aclSource attribute 253	directory hierarchy 4
adding to directory 290	DNs 156
arranging 4	files shipped with LDAP server 39
creating for referrals 277	glddebug 99
data model 155	JCL for LDAP server 413
defining 155	JCL for Idif2tdbm 415
description 4	JCL for tdbm2ldif 417
entryOwner attribute 253	LDAP URL 60
identifying 155	LDIF file 291
loading 7, 290	LDIF input file 421
loading from TDBM 137	modifying user in RACF 220
loading into TDBM 127	object class hierarchy 165
loading large number of 127	overrides 257
ownerPropagate attribute 254	partitioning DB2 tables for TDBM 42
ownerSource attribute 254	permissions 254
permissions 251	propagation 257
protecting 249	referral objects 278
TDBM directory schema 159	referrals, distributing namespace 280
entryOwner attribute	removing user from group in RACF 222
description 253	removing user from RACF 222
migrating 106	replica object definition 272
environment variables	samples 39 schema entry 160
LANG 151	searching for user information in RACF 220
LDAP SLAPD ENVVARS FILE 125	scarcing for user information in RACE 220

examples (continued)	finding
searching for user's connection to group 222	subschemaSubentry DN 176
setting up Kerberos directory 229	first time installing 12
setting up native authentication 238	formats
slapd.conf 371	GIF 4
slapd.envvars file 151	JPEG 4
SPUFI script (TDBM) 403, 409	slapd.conf 52
using ref attribute 277	front end
using schema2ldif utility 180	for X.500 6
EXOP backend	
configuring 32	
configuring, example 84	G
GetDnForUserid 247, 435	gathering group memberships 57
GetPrivileges 247, 436	gathering trace records 98
migration information 119	GDS (Global Directory Service)
PC callable support mode 76, 79	See Global Directory Service (GDS) 6
section in slapd.conf 51	Generalized Time syntax 175
setting up for 43	GetDnForUserid 247, 435
using 247	GetPrivileges 247, 436
using for Policy Director 8	GIF format 4
extattr command 37	glddebug 99
extended group membership searching 9, 57	global configuration file options 53
extended operations	Global Directory Service (GDS) 6
messages 367	global section in slapd.conf 51
extended operations backend	glossary of terms 455
See EXOP backend 8	group membership 432
extendedGroupSearching option 57	group name 71
extensibleObject object class 172	groups
external bind, SASL 8	access control 258
external security manager 35	connecting to in RACF 215
	creating for ACL 266
F	delete for ACL 268
	extended, membership searching 9, 57
files	modify for ACL 268
configuration 371	universal, searching 218
configuration, global options 53	GSS API bind 225, 227
copying configuration 39	
DB2 CLI initialization 14	Н
envvars 151	
generating 291	Hardware Configuration Definition (HCD)
Idap.db2.profile file 27	configuring for 11
ldap.msys.db2.profile 29 ldap.msys.profile 29	hash formats 48
	HCD (Hardware Configuration Definition)
ldap.msys.racf.profile 29 ldap.msys.slapd.profile 29	See Hardware Configuration Definition (HCD) 11
	HFS (Hierarchical File System)
ldap.profile 20 ldap.racf.profile file 27	See Hierarchical File System (HFS) 6
Idap.slapd.profile file 27	Hierarchical File System (HFS)
LDIF format 290	changing debug setting 98
profile 17	files, converting 97
sample LDIF input 421	naming the LDAP server 6
schema.IBM.Idif 176	hierarchical tree
schema.user.ldif 176	defining 4
schema.user.idii 176 slapd.conf 52, 371	hierarchy
SPUFI 403	directory 155
stash 70	directory, example of 4
TDBMDB SPUFI 403	example, object class 165
TDBMINDX SPUFI 409	laying out entries in 288
filters	referrals 278
searching 256	
Joan Jilling 200	

	J
i 16	Japanese messages 151
IBM attribute types	JCL (Job Control Language)
description 164	See Job Control Language (JCL) 93
format 171	JOB card, modifying 125
usage 171	Job Control Language (JCL)
IBM-1047 character set 54	for running SLAPD as started task 93
ibm-entryuuid	generated by Idapcnf 17
replication 269	running DB2 backend utilities from 125
ibm-entryuuid,attribute 9	sample for LDAP server 413
IBMAttributeTypes schema attribute 164, 171	sample for Idif2tdbm 415
ICSF (Integrated Cryptographic Services Facility)	sample for tdbm2ldif 417
See Integrated Cryptographic Services Facility	JPEG format 4
(ICSF) 16	
identity mapping	1.7
Kerberos 226	K
identity, Kerberos 36	KDC (Key Distribution Center) 225
idleConnectionTimeout option 57	Kerberos
include option 58	authentication
index option 58	description 225
indexes, create DB2 40, 409	bind capability 8
indexing	configuring access control 227
specifying type of 58	configuring administrator for 81
information	identity
arranging 4	defining 36
layout 287	identity mapping 226
protecting 6	identity mapping option 58
referencing 5	installing 16
inheritance	migration information 117
default 254	operating environments 231
initialization file, DB2 CLI 14	setting up for 225
initializing	settup up directory example 229
native authentication 233	specifying administrator identity 59
replica database 272	specifying attribute in SDBM 215
Installation Verification Procedure (IVP)	specifying identity 67
for configuring LDAP server 27	specifying keytab 58
installing PR2 42	specifying participation in 71
DB2 13	updating directory schema 226
first time 12	using Idapenf with 28
for CLI 13 for ODBC 13	key database file
	specifying 73
ICSF for password encryption 16 LDAP server 11	specifying for server 69
migrating 103	Key Distribution Center (KDC) 225
OCSF for password encryption 16	key label 65
Policy Director 15	keyboard 447
RACF 15	keytab file
related products 13	generating for server 225
roadmap 11	krbldentityMap option 58
System SSL 15	krbKeytab option 58 krbLDAPAdmin option 59, 81
Integrated Cryptographic Service Facility (ICSF)	KIDLDAPAdifilif option 59, 61
installing for password encryption 16	
interactions, backend variables 33	I
interface changes 123	
international characters	IA5 character set 298
retrieving 8	LANG environment variable 151
storing 8	LANG parameter 35
IVP (Installation Verification Procedure)	language setting 35
See Installation Verification Procedure (IVP) 27	large number of entries, loading 127 layout, information 287
,	iayuu, iiiiuiiiailuii Zo <i>i</i>

LDAP (Lightweight Directory Access Protocol)	LDAP server (continued)
See Lightweight Directory Access Protocol	model for 4
(LDAP) 4	multi-server mode 75, 76
LDAP administrator	multiple single-server mode LDAP servers 75
role for Idapcnf 23	multiple, setting up 79
LDAP client	name, installation 35
authenticateOnly server control 432	naming 6
cipher specifications 68	new information xxi, xxii
considerations 297	NLS 151
requesting number of threads 73	object classes supported 171
using in LDAP 6	planning for 11
UTF-8 data 298	preparing 35
LDAP client requests	process ID 97
listening for 59 LDAP configuration utility 17	RACF 205 RDBM collation 151
LDAP Data Interchange Format (LDIF)	replica 272
description 290	replication 269
input file, sample 421	requirements for user ID 35
loading entries from TDBM 137	retrieving ACL information 258
loading entries into TDBM 127	retrieving Policy Director data 8
LDAP directory	running as started task 93
protecting information in 6	sample 12, 371
LDAP directory server	sample files 39
See LDAP server 3	sample JCL 413
LDAP schema attributes 167	SDBM backend 205
LDAP server	setting up SDBM 42
access control 249	shutting down 98
administrator 54	single-server mode 75
administrator DN 80	started task, defining 93
alternate server 55	starting from console 94
attribute types supported 170	starting in SDSF 93
authenticateOnly server control 432	starting in shell 97
capabilities 6	starting up 13
changed information xxiii	stopping 97
changing replica to master 274	stopping from console 96
command-line options 94	stopping in SDSF 96
configuration options 51	table spaces for TDBM 40
configuring 75	TDBM collation 151
configuring for multi-server mode 77	TDBM database, creating 40
configuring with Idapcnf utility 17	tracing 98
creating user ID for 35	user ID 33
data model 155	using 6
DB2 server location 67	using for user ID 35
debugging facility 98	verifying 102
default directory 35	Version 3 protocol 8
defining separate user ID for 97	LDAP syntaxes
defining started task for 93	description 166
defining user ID 35	format 167
description 3	supported, general use 168
enabling SSL for 46	supported, server use 168
envvars file 151	usage 167
equivalent server 55	LDAP URL
example configuration 288	examples 60
extended group membership searching 9, 57	starting LDAP server 94
IBM attribute types supported 171	LDAP URL format
installing 11, 13	specifying for listen 59
LDAP syntaxes supported 167 master and replica 274	LDAP_DEBUG environment variable 95
matching rules supported 169	LDAP_SLAPD_ENVVARS_FILE setting 97
messages 317	Idap_ssl_client_init API
migrating 103	using for SASL bind 47
migrating 103	using ior sase billio 41

Idap_ssl_init API	levels, debug 94
using for SASL bind 47	licensed publications xviii
ldap.db2.profile file 27	Lightweight Directory Access Protocol (LDAP)
ldap.msys.db2.profile 29	description 4
Idap.msys.profile 29	how it works 6
Idap.msys.racf.profile 29	schema publication 159
Idap.msys.slapd.profile 29	Lightweight Directory Access Protocol (LDAP) server
Idap.profile file 20	See LDAP server 3
Idap.racf.profile file 27	limit, time
Idap.slapd.profile file 27	specifying in configuration file 72
Idapadd utility	limitations, referrals 279
loading entries into TDBM 131	listen option 59, 73, 79, 106
Idapadduuids utility	loading
running from JCL 141	directory information 290
running in z/OS shell 141	entries 7
Idapcnf utility	entries from TDBM 137
advanced options 27	entries into TDBM 127
capabilities 18	large number of entries 127
•	locating
configuration confirmation 27	slapd.conf 51
input files 20	subschemaSubentry DN 176
messages 366	logfile option 61
msys use of 29	LookAt messages xviii
overview 17	LOOKAT Messages XVIII
restrictions 18	
	M
steps for configuring with 24	
using 20	mail attribute 4
Idapdelete	Managed System Infrastructure (msys)
deleting group entry 268	LDAP configuration for 29
Idapmodify	manageDsalT server control 280, 431, 432
updating ACLs 259	mapping
LDAPResult construct 299	identity, Kerberos 226
Idapsearch utility	Kerberos default 227
using to verify LDAP server 102	LDAP-style names to RACF attributes 206
LDAPSRV PROC JCL 413	mask for debug level 94
LDAPSRV user ID 35, 93	master
LDIF (LDAP Data Interchange Format)	backup of 270
See LDAP Data Interchange Format (LDIF) 127	changing replica to 274
LDIF files	communicating with replica 274
editing 291	database, description 269
Idif2db utility	server 269, 270
encrypting user passwords 48	server DN 61
initial loading of database 290	server password 62
running from JCL 125	server, setting up 275
running in shell 125	server, specifying 61
running in TSO 126	using replication 293
using with replica server 274	masterServer option 61
Idif2tdbm utility	masterServerDN option 61
allocating datasets for 129	masterServerPW option 62
description 127	matching rules
encrypting user passwords 48	description 166
entering ACLs 249	EQUALITY values 163
initial loading of database 290	EQUALITY, defaults 162
loading adminDN entry 80	format 169
recovery 133	ORDERING values 163
running from JCL 125	SUBSTR values 163
running in shell 125	supported 169
running in TSO 126	usage 169
sample JCL for 415	maxConnections option 62, 108
TDBM schema 174	maxThreads option 56, 62, 73, 107
less than sign 156	MAY attribute type 165, 172

MD5 encryption format 48 MD5 encryption method 64	multiple LDAP servers setting up 79
mdoify DN	multiple socket ports 9
access control 189	multiserver option 63
membership	MUST attribute type 165, 172
extended group, searching 9, 57	mutual authentication 8
messages extended operations 367	
LDAP server 317	N
Idaponf 366	
SDBM 365	namespace
started task 96	directory 287
TDBM 345	entries, RACF 210 hierarchy diagram for RACF 210
using LookAt xviii	National Language Support (NLS)
migrating	setting variables for 151
actions required for 104	native authentication
from previous LDAP releases 103	capability 8
RDBM to TDBM 176	defining participation 233
roadmap 110	description 233
schema2ldif utility 177	enabling 72
schemas 177	example of setting up 238
steps for 103	initializing 233
TDBM databases 108	installing RACF for 15
migrating TDBM databases 108	migration information 118
migrating TDBM schema 439	operating modes 234
minimum schema for TDBM 393	password changes 63
modes	specifying DN 63
choosing multiserver 63	updating schema 233
multi-server 75, 76	using with Web servers 241
multi-server, operating in 77	native operations
PC callable support 76, 79 single-server mode 75	running 237
single-server mode LDAP servers 75	nativeAuthSubtree option 63, 234
single-server, operating in 76	nativeUpdateAllowed option 63, 234
modify DN	Network Authentication and Privacy Service 225 new information xxi, xxii
access control changes 191	NLSPATH environment variable 151
activity log 203	normal access class 251
complex replication 269	nstalling
considerations 188	Kerberos 16
considerations in the use of 187	numeric object identifier (OID) 172
entries for rename 188	, , ,
operation syntax 183	
operations and replication 200	0
relocating an entry 190	object class
replication synchronization 201	adding for TDBM 171
scenario constraints 194	definitions, adding for TDBM 171
SDBM schema 211	description 155, 165
suffix DNs 193	extensibleObject 172
Modify DN validation 201	files, converting 177
msys (Managed System Infrastructure)	format for TDBM 171
See Managed System Infrastructure (msys) 29	hierarchy 165
multi-server mode	Kerberos 226
configuring example 85	person 287
configuring for 77	referral 277
running in 75, 76	replicaObject 270
multi-valued ref attribute 277	TDBM directory schema 160
multiple databases	TDBM usage 171 object class attribute
concurrent 7	description 5
replication of 269	object class definitions
	adding for TDBM 171

object class definitions (continued)	password (continued)
specifying for TDBM 171	master server 62
objectClass attribute 160	native modify 237
objects	RACF
protecting 249	changing using SDBM 219
referral 278	replication key database 73
replica 270	SSL key database file 70
replica, adding in TDBM 272	storing in stash file 70
OCSF (Open Cryptographic Services Facility)	password encryption
See Open Cryptographic Services Facility	configuring for 32, 47
(OCSF) 16	db2pwden utility 148
ODBC (Open Database Connectivity)	description 9
See Open Database Connectivity (ODBC) 13	installing ICSF for 16
oedit editor 291	installing OCSF for 16
OGET command 97	protecting environment for 37
OID (numeric object identifier) 172	pwEncryption configuration option 64
one-way hash formats 48	replication 269
Open Cryptographic Services Facility (OCSF)	unloading TDBM database 137
installing for password encryption 16	PC (program call)
Open Database Connectivity (ODBC)	See program call (PC) 64
setting up for 13	pcThreads option 64
operating environmnets	PDS (partitioned data set)
Kerberos 231	See partitioned data set (PDS) 93
operating modes	performance improvements 7
native authentication 234	performance information
operations	capturing 100
defining 6	performance, LDAP server 7
operator console	periodic checks
starting SLAPD from 94	modify DN 201
options	permissions
checklist 52	access 250
configuration file 53	attribute access classes 252
organization of book xvii	determining 254
organizing	entry 251
directory namespace 287	examples 254
information 4	places, modeling information for 287
out-of-sync conditions 275	plan name 33
overhead, reducing 7	PLANNAME value 33
override, ACL example 257	planning
owner	LDAP server setup 11
creating for entry 263	plus sign 156
deleting for entry 266	Policy Director
modifying for entry 264	data, accessing 247
ownerPropagate attribute 254	data, retrieving for LDAP 8
ownerSource attribute 254	migration information 119
ownerPropagate attribute	retrieve DNs 435
description 254	retrieving data 247, 436
ownerSource attribute	setting up to support 15
description 254	setting up to use 43
	populating
D	replica database 272
P	port
Parallel Sysplex	multiple socket 9
configuraiton example 85	numbers, default 94
server name 72	TCP/IP for SSL 66
using 75, 76	TCP/IP, specifying 64
partitioned data set (PDS)	port option 64, 73, 106
DLLs 93	pound sign (#) support in SDBM 211
password	pre-configuration set up 13
administrator, specifying 55	preparing
changing 63	for backend interactions 33

preparing (continued)	RDBM backend (continued)
LDAP server 35	collation 151
printing	configuring 32
trace tables 99	installing and setting up 13
procedure	managing administrator DN 80
JCL to run SLAPD 93	managing password 80
process ID	migrating schema to TDBM 176
providing for SLAPD 97	migrating to TDBM 106
profile files	password encryption 9, 47, 64
used with Idapcnf 17	preparing for 33
used with Idaponf for msys 29	running utilities for 125, 126
program call (PC)	using 7
callable support mode 76, 79	using password encryption 147
initializing threads for 64	verifying 102
using LDAP server for callable support 102	RDN (relative distinguished name)
program control	See relative distinguished name (RDN) 5
for LDAP server 37	
	readOnly option 65
for SSL 37	reason codes 299
OCSF DLLs for password encryption 37	recovering
propagation, ACL	from out-of-sync conditions 275
description 256	ldif2tdbm 133
example 257	ref attribute 277
indicating, flag for 253	references
protecting	setup, recommended 278
environement for password encryption 37	referencing information 5
environment for LDAP server 37	referral option 65
environment for SSL 37	referrals
information 6	default 65, 290
information using ACLs 249	default, defining 278
protecting access 70, 73	description 277, 293
protection, scope of	example of distributing namespace 280
attribute privileges 250	LDAP server capability 7
determining 250	limitations with version 2 279
protocol	manageDsalT server control 431, 432
directory 4	object 293
Version 3 8	processing 279
PROXY segment 36	replication 274
•	
pseudo DN 252	specifying 65, 290
publication, schema 159	suppressing 431, 432
publications, related 453	using with dynamic WLM 78
pwEncryption option 64, 148, 269	using without dynamic WLM 77
	Version 2 protocol 279
0	Version 3 protocol 280
Q	relational database
query command 100	loading entries into TDBM 127
querying	relative distinguished name (RDN)
root DSE 297	description 5, 155
schema 160	Release 1 (z/OS LDAP)
subschemaSubentry 161	new and changed functions 121
quotation marks	update summary 111
using in configuration file 52	Release 10 (OS/390 LDAP)
doing in configuration inc - 52	changes in book xxiii
	new and changed functions 121
R	update summary 111
	Release 2 (z/OS LDAP)
R_Admin interface 216	changes in book xxii
RACF (Resource Access Control Facility)	now and abanged functions 117
See Resource Access Control Facility (RACF)	55
RACF key rings	update summary 110
using 145	Release 4 (z/OS LDAP)
RDBM backend	new and changed functions 111
buffer pools 14	update summary 110

relocating an entry with DN realignment 190 removing schema file include statements 110 removing verifySchema configuration option 110	roadmap configuring LDAP server 31 installing and running LDAP server 11
replica	migration 110
changing to master 274	roles
communicating with master 274	configuration utility 22
master server 61	root DSE
objects 275	searching 297
setting up 274, 293	support of 9
replication	RRSAF (Recoverable Resource Manager Services
associating servers with 280	Attachment Facility)
benefits 270	See Recoverable Resource Manager Services
database,description 269	Attachment Facility (RRSAF) 15
description 269	rules, matching
ibm-entryuuid 269	See matching rules 169
LDAP server capability 7	running
objects 270	DB2 backend utilities 125, 126
objects, adding in TDBM 272	LDAP server as started task 93
password encryption 269	LDAP server in z/OS shell 97
RDBM 75	LDAP server using data sets 97
server 272	LDAP server using DB2 13
server, configuring 273	LDAP tools with SDBM 212
setting up for 293	native operations 237
specifying key database file for 73	password encryption utility 147
SSL 274	roadmap 11
TDBM 75	
troubleshooting 275	
replKeyRingFile option 73	S
replKeyRingPW option 73	SAF (Security Authorization Facility) 15
Request for Comments (RFC)	samples
supported by z/OS LDAP 9	DSNAOINI file 15
requested attributes 256	JCL for LDAP server 413
requests, client	JCL for tdbm2ldif 417
processing 62	LDAP server 12
resetting	LDIF input file 421
trace table 99	object class hierarchy 165
Resource Access Control Facility (RACF)	schema entry 160
accessing information in 42	slapd.conf 371
administrator DN 81	slapd.envvars file 151
changing password with SDBM 219	SPUFI script (TDBM) 403, 409
command to create LDAPSRV 36	using schema2ldif utility 180
commands for defining started task 93	SASL CRAM-MD5 and DIGEST-MD5 8
configuring LDAP server for 8	SASL external bind 8
connection to group 215	SASL GSS API Kerberos bind 8
distinguished names 156	SCEERUN dataset 93
input file for Idapcnf 27	scenarios
installing for native authentication 15	configuration 82
installing for SDBM 15	schema
LDAP access to 205	See also schema (TDBM)
mapping attributes 206	dynamic 8
namespace entries 210	migrating 108
password 81	modifying 127
PROXY segment 36	updating for Kerberos 226
restriction on amount of output 219	updating for native authentication 233
Subsystem function 15	verifying 73
univeral groups	schema (RDBM)
searching 218	replica object 270
using 35	schema (TDBM)
RFC (Request for Comments)	attribute syntax 166
See Request for Comments (RFC) 9	attribute types 170
	defining new elements 172

schema (TDBM) (continued)	SDBM schema (continued)
directory 159	modify DN 211
entry, displaying 175	SDSF
errors, analyzing 174	file 96
IBM attribute types 171	starting LDAP server in 93
LDAP attributes 167	searching
LDAP syntaxes 167	across multiple servers 277
matchng rules 169	anonymous 256
migrating RDBM to TDBM 176	directories 6
minimum schema 393	entire RACF database 218
object classes 171	permissions required 256
publication 121, 159	replication 270
replica object 270	root DSE 297
retrieving TDBM 175	schema entry 176
sample entry 160	using attributes 256
searching for schema entry 176	Secure Sockets Layer (SSL)
update 122, 159	certificate 68
updating TDBM 173	certificate authentication 68
schema file	cipher specifications 68
removing include statements 110	configure LDAP server using 7
schema-updates.ldif 160	configuring 32
schema.IBM.ldif file 176	enablement 274
schema.user.ldif file 176	enabling 46
schema2ldif utility 177	key database file
scope of protection	protecting access to 70, 73
attribute privileges 250	specifying 73
determining 250	specifying for server 69
scripts	password for key database file 70
modifying for TDBM 41	preparing for 15
running for TDBM 41	protecting environment for 37
SDBM backend	replication 274
changing password in RACF 219	setting up options for 45
configuration utility 18	specifying in configuration file 66
configuring for 31	stash file 70
connection to group 215	TCP/IP port for 66
deleting attributes 222	securePort option 66, 73, 106
implementing 8, 205	SecureWay Security Server Network Authentication and
installing RACF for 15	Privacy Service for z/OS 16
Idapcnf utility 18	security
mapping, Kerberos 227	options, setting up 45
messages 365	specifying type of 66
migration information 118	security administrator
namespace hierarchy 210	role for Idaponf 23
pound sign (#) support 211	Security Authorization Facility (SAF)
preparing for 33	using for Policy Director support 15
RACF-style DNs 156	security manager, external 35
running LDAP tools with 212	security option 46, 66, 73, 106
running with other backends 43	Security Server for z/OS 3
running with TDBM 43	See Security Authorization Facility (SAF) 15
running without RDBM 43	segment, PROXY 36
running without TDBM 43	sendV3stringsoverV2as option 66, 298
schema publication 159	sensitive access class 251
section in slapd.conf 51	server
setting BINDPW in PROXY segment 36	alternate 55
setting up for 42	associating with referrals 278
setting up user ID 36	certificate 46
support of PROXY segment 36	master 270, 275
using for authentication 249	master, problems 275
using LDAP operation utilities with 220	master, specifying 61
verifying 102	name, specifying 72
SDBM schema 397	parent 278
ODDINI SOLICITIC 131	parent 210

server (continued)	SSL/TLS
pointing to others 278	creating key database or key ring 45
referrals 277	enabling 44
replica 272, 274	LDAP utilities 145
using in LDAP 6	obtaining a certificate 45
server controls	protected communications 44
authenticateOnly 57, 432	setting up an LDAP client 47
LDAP server capability 9	setting up for 43
manageDsalT 280, 431, 432	sslAuth option 47, 68
server name 33	sslCertificate option 68
server replication	sslCipherSpecs option 68
modify DN 201	sslKeyRingFile option 69, 73
serverEtherAddr option 66	sslKeyRingFilePW option 70
serverKrbPrinc option 67	sslKeyRingFilePWStashFile option 70
servername option 67	sslKeyRingPWStashFile option 73
ServerPac installation 93	stand-alone LDAP daemon
setting up	See LDAP server 4
buffer pools 14	start TLS 437
DB2 13	started task
directory namespace 287	changing debug setting 98
for CLI 13	defining 93
for ODBC 13	defining 93 defining user ID for 35
Kerberos directory 229	messages 96
•	
LDAP server using roadmap 11	running LDAP server as 93
multiple LDAP servers 79 native authentication example 238	starting LDAP server in SDSF 93
•	
related products 13	loading DLLs 93
TEMP datasets 14	SLAPD from console 94
TEMP space 14	stash file 70
SHA encryption format 48	steps for installing and running LDAP 11
SHA encryption method 64	steps for migrating 103
shortcut keys 447	stopping
shutting down	LDAP server from console 96
LDAP server 96, 98	LDAP server in SDSF 96
single-server mode	SLAPD 97
configuring for 76	storing information 4
multiple 75	STRUCTURAL object class type 165, 172
replicating in 269	subject
running in 75	determining rights for 250
sizelimit	submit command 125
specifying in configuration file 67	subschemaSubentry attribute 161
sizeLimit option 67	subschemaSubentry DN 176
SLAPD	suffix
See LDAP server 3	option 70, 80
slapd arguments 107	summary of changes xxi, xxii, xxiii
slapd.conf	supportKrb5 option 71
creating 51	synchronizing databases 269, 272, 294
format 51	syntax
locating 51	DN 156
See configuration file 52	EQUALITY matching rules 162, 163
slapd.envvars file 151	ORDERING matching rules 163
socket ports, multiple 9	schema attribute 166
space, white 156	SUBSTR matching rules 163
spaces, blank 52, 156	sysplex
SPUFI (SQL Processor Using File Input) facility	See Parallel Sysplex 75
See SQL Processor Using File Input (SPUFI)	sysplexGroupName option 71
facility 40	sysplexServerName option 72
SQL Processor Using File Input (SPUFI) facility	system administrator
creating TDBM database with 40	role for Idapcnf 23
TDBMDB SPUFI script (TDBM) 403	SYSTEM file, creating 131
TDBMINDX SPUFI script (TDBM) 409	

system programming	TDBM backend (continued)
role for Idapcnf 23	setting up multiple servers 79
	SPUFI 403
_	SPUFI script 409
T	tdbm2ldif utility 137
table	using 7
in-storage trace 99	using password encryption 147
table spaces for TDBM	values for SPUFI 40
creating 40	verifying 102
partitioning 41	TDBM schema
task, started	setting up 159
See started task 93	updates between releases 160
TCP/IP (Transaction Control Protocol/Internet Protocol)	upgrading 159
	TDBM schema migration 439
See Transaction Control Protocol/Internet Protocol	tdbm2ldif utility
(TCP/IP) 4	description 137
TDBM backend	encrypting user passwords 48
See also schema (TDBM)	replicating 272
access control	running from JCL 125
pseudo DN 252	
ACL attributes 249	running in shell TDBM backend
adding replica objects 272	running utilities for 125
attribute types supported 170	running in TSO 126
buffer pools 14	sample JCL for 417
collation 151	TDBMDB SPUFI file 403
configuration utility 18	TDBMINDX SPUFI file 409
configuring 32	TEMP datasets
configuring administrator DN 80	setting up 14
configuring password 80	terms, glossary of 455
configuring sample server with 39	threads
configuring with multi-server 85	for performance 7
configuring with SDBM 83	specifying number in configuration 64
configuring with SSL and password encryption 82	specifying number of 62, 73
CRAM-MD5 and DIGEST-MD5 243	specifying with commThreads 55
creating DB2 databases 40	ticket
creating table spaces 40	Kerberos authentication 225
directory schema 159	Time Sharing Option (TSO)
dynamic schema 8, 159	running backend utilities in 126
entry owner attributes 249	timeLimit option 72
IBM attribute types supported 171	timeout
initializing ACLs 254	specfiying 57
installing and setting up 13	timestamp changes 106
LDAP syntaxes supported 167	TLS
Idapcnf utility 18	start 437
ldif2tdbm utility 127	tracing
mapping, Kerberos 227	options 98
matching rules supported 169	using in-storage trace table 99
messages 345	Transaction Control Protocol/Internet Protocol (TCP/IP)
migrating from RDBM 106	port for SSL 66
migrating schema from RDBM 176	port, specifying 64
migration information 121	running over 4
native authentication 72, 233	tree structure
object class supported 171	hierarchical 4
partitioning DB2 tables 41	troubleshooting
password encryption 47	reason codes 299
	replication 275
preparing for 33	TSO (Time Sharing Option)
RDN 155	See Time Sharing Option (TSO) 126
replica server 272	two-way encryption format 49
running utilities for 125, 126	types of information to store 4
running with SDBM 43	types of illionnation to stole 4
schema attribute syntax 166	
section in slapd.conf 51	

U	Version 3 protocol
UID O	LDAP support of 8
running under 94	LDIF files 138
Unicode	referrals 280
See UTF-8 characters 61	TDBM schema 105, 159
universal groups	UTF-8 characters 151, 298
searching 218	
unsynchronized databases 275	W
update, TDBM schema 159, 173	- -
updating	waitingThreads option 56, 73, 107
schema errors 174	Web servers using native authentication with 241
schema for native authentication 233	WebSphere Application Server
useNativeAuth option 63, 72	configuring for 11
user ID	white space 156
creating LDAP server 35	mino opaco Too
defining for LDAP server 35	
defining separate 97	X
specifying for owner 56 user password encryption	X.500
See password encryption 9	data model 155
user-defined schema 176	description 6
userNativeAuth option 234	X.509 standard
userPassword attribute value	digital certificate 46
See also password encryption	_
encrypting 148	_
specifying encryption method for 64	Z
unloading encrypted 137	z/OS Security Server 3
UTC Time syntax 175	z/OS shell
UTF-8 characters	running LDAP server in 97
mapping with Unicode 151	
retrieving 8	
storing 8, 298	
using in DNs 54, 61	
utilities	
db2pwden 148	
LDAP operation	
using with SDBM 220	
ldapcnf 17 Idif2tdbm 127	
schema2ldif 177	
tdbm2ldif 137	
tabinizian 107	
V	
V2 protocol	
See Version 2 protocol 279	
V3 protocol	
See Version 3 protocol 8	
validateincomingV2strings option 73	
variable interactions, backend 33	
verifying	
LDAP server 102	
verifying configuration 27	
verifySchema option 73	
Version 2 protocol	
output data format 66	
referrals 279	
referrals, limitations 279	
validating data 73	

Readers' Comments — We'd Like to Hear from You

z/OS Security Server LDAP Server Administration and Use

Publication No. SC24-59	923-03				
Overall, how satisfied are you with the information in this book?					
Overall satisfaction	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
How satisfied are you th	at the information	in this book is:			
Accurate Complete Easy to find Easy to understand Well organized Applicable to your tasks	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Please tell us how we ca	an improve this bo	ok:			
Thank you for your respon	nses. May we conta	ct you?	s 🗌 No		
When you send comment way it believes appropriate				r distribute your c	omments in any
Name		Ad	dress		
Company or Organization					
Phone No.					

Readers' Comments — We'd Like to Hear from You SC24-5923-03



Cut or Fold Along Line

Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation Department 55JA, Mail Station P384 2455 South Road Poughkeepsie, NY 12601-5400



laalladallallaanalladadalllaallaanalll

Fold and Tape

Please do not staple

Fold and Tape

IBM.

Program Number: 5694-A01, 5655-G52

Printed in U.S.A.

SC24-5923-03

